Space Bounds for Adaptive Renaming

Maryam Helmi[1], Lisa Higham[2], and Philipp Woelfel[3]

University of Calgary
Department of Computer Science
Calgary, T2N1N4 Alberta, Canada

arXiv:1603.04067v6 [cs.DC] 15 May 2016

Abstract. We study the space complexity of implementing long-lived and one-shot adaptive renaming from multi-reader multi-writer registers, in an asynchronous distributed system with n processes. As a result of an *f-adaptive renaming algorithm* each participating process gets a distinct name in the range $\{1,\ldots,f(k)\}$ provided k processes participate.

Let $f:\{1,\ldots,n\}\to\mathbb{N}$ be a non-decreasing function satisfying $f(1)\leq n-1$ and let $d=\max\{x \mid f(x)\leq n-1\}$. We show that any non-deterministic solo-terminating long-lived f-adaptive renaming object requires $d+1$ registers. This implies a lower bound of $n-c$ registers for long-lived $(k+c)$-adaptive renaming, which we observe is tight.

We also prove a lower bound of $\lfloor\frac{2(n-c)}{c+2}\rfloor$ registers for implementing any non-deterministic solo-terminating one-shot $(k+c)$-adaptive renaming. We provide two one-shot renaming algorithms: a wait-free algorithm and an obstruction-free algorithm. Each algorithm employs a parameter to depict the tradeoff between space and adaptivity. When these parameters are chosen appropriately, this results in a wait-free one-shot $(\frac{3k^2}{2})$-adaptive renaming algorithm from $\lceil\sqrt{n}\rceil+1$ registers, and an obstruction-free one-shot f-adaptive renaming algorithm from only $\min\{n,x \mid f(x)\geq 2n\}+1$ registers.

[1] mhelmikh@ucalgary.ca, +1 403 210-9416
[2] higham@ucalgary.ca, +1 403 220-7696
[3] woelfel@ucalgary.ca, +1 403 220-7259

1 Introduction

Distributed systems with a large number of processes, such as the Internet, provide services that are typically used by only a small number of processes simultaneously. This is problematic if the time or space used by the service is a function of the size of the name-space of the processes that could use it. The time or space consumed by such applications can be significantly decreased by having each process that wants to use the application first acquire a temporary name from a name space that is adequate to distinguish all the participants, but much smaller than the name-space of the distributed system, and then return the temporary name to the pool when it is finished with the service. This is the role of a shared renaming object. A related application of the renaming object is in operating systems where processes repeatedly acquire and release names that correspond to a limited number of resources [12]. Renaming is an important tool in distributed computing [5] because it enhances the practicality and usefulness of distributed system services. A renaming object may be even more useful if the time and space resources it consumes is a reasonable function of the actual number of processes that are currently either holding, acquiring, or releasing a name. In this paper, we address the renaming problem for the standard asynchronous shared memory model with n processes.

With an f-adaptive renaming object, each of the n processes can perform a `getName()` and return a distinct name in a small domain $\{1,\ldots,f(k)\}$ where k is the number of participants. Herlihy and Shavit [24], and also Rajbaum and Castañeda [13] showed that there is no deterministic, wait-free implementation of $(2k-2)$-adaptive renaming from multi-reader multi-writer registers. This result also follows from the relationship between the adaptive renaming problem and *strong symmetry breaking* (SSB): a $(2k-2)$-renaming algorithm implies a solution to SSB [7], for which there is no deterministic wait-free solution [7, 9, 20]. This impossibility can be circumvented by using randomness or stronger primitives such as compare-and-swaps [2, 4, 16, 28, 29]. The step complexity of deterministic and randomized algorithms has been studied extensively in asynchronous systems (see e.g., [1, 3, 4, 10, 17]). However, there are no previous results on the space complexity of adaptive renaming. Because renaming seems to require that participants discover information about each other, adaptive renaming appears related to f-adaptive collect. A collect algorithm is f-adaptive to total contention, if its step complexity is $f(k)$, where k is the number of processes that ever took a step. Attiya, Fich and Kaplan [6], proved that $\Omega(f^{-1}(n))$ multi-reader multi-writer registers are required for f-adaptive collect.

Suppose you have m shared registers available to construct a renaming object for a system with n processes. First we would like to know under what additional conditions such an implementation exists, and when it does, how best to use the m registers. Suppose, when there are k participants, the acquired names are in the range $\{1,\ldots,f(k)\}$. Will $f(k)=k^c$ for a small constant c suffice for the application? Must $f(k)$ be closer to k, say within a constant? Perhaps it should even be exactly k (*tight* adaptive renaming)? Does the application need to permit processes to repeatedly acquire and release a name (*long-lived* renaming), or do processes get a name at most once (*one-shot* renaming)? How strong a progress guarantee is required? Is the number of participants usually less than some bound b much smaller than n? If so, is there still some significant likelihood that the number of participants is somewhat bigger than b, or is there confidence that the bound b is never, or only very rarely, exceeded? In the rare cases when there are a large number of participants, can the system tolerate name assignments from a very large name space?

In order to study the space complexity implication for these questions, we first generalize the adaptive renaming definition. Both versions (*long-lived* and *one-shot*) of *b-bounded f-adaptive renaming* support the operation `getName()`, which returns a name to each invoking process. The long-lived version also supports an operation `relName()`, which releases the name to the available name domain. Both versions must satisfy 1) no two processes that have completed a `getName()` and have not started their following `relName()`, receive the same name, 2) if there are $k \le b$ processes that have invoked `getName()` and have not completed their subsequent `relName()` during an execution of `getName()` by process p, then p returns a name in $\{1,\ldots,f(k)\}$. Observe that for the one-shot case, k is the number of processes that have started a `getName()` before p completes its `getName()`. We call the problem of n-bounded f-adaptive renaming simply f-adaptive renaming. The special case when $f(k)=k$ and $b=n$ is called *tight* renaming. Our goal is to determine the relationships between b, $f(k)$, and m for one-shot versus long-lived, and wait-free versus non-deterministic solo-terminating implementations of adaptive renaming objects from multi-reader/multi-writer registers.

Let $f:\{1,\ldots,n\}\to\mathbb{N}$ be a non-decreasing function satisfying $f(1)\le n-1$ and let $d=\max\{x \mid f(x)\le n-1\}$. Note that if $f(1)\ge n$, f-adaptive renaming is a trivial problem. In this paper we show:

- At least $d+1$ registers are required to implement any non-deterministic solo-terminating long-lived d-bounded f-adaptive renaming object.
- At least $\lfloor \frac{2(n-c)}{c+2} \rfloor$ registers are required to implement any non-deterministic solo-terminating one-shot $(k+c)$-adaptive renaming object where, c is any non-negative integer constant.
- For any $b \leq n$, there is a wait-free one-shot $(b-1)$-bounded $(k(k+1)/2)$-adaptive renaming algorithm implemented from b bounded registers. When $k \geq b$, the returned names are in the range $\{1,\ldots,n+\frac{b(b-1)}{2}\}$.
- For any $b \leq n$, there is an obstruction-free one-shot $(b-1)$-bounded k-adaptive renaming algorithm implemented from $b+1$ bounded registers. When $k \geq b$, the returned names are in the range $\{1,\ldots,n+b-1\}$.

By using these results and setting b appropriately we then derive the following corollaries:

- A wait-free one-shot $(\frac{3k^2}{2})$-adaptive renaming algorithm that uses only $\lceil\sqrt{n}\rceil+1$ registers.
- An obstruction-free one-shot f-adaptive renaming algorithm that uses only $\min\{n,x \mid f(x) \geq 2n\}+1$ registers.
- A tight space lower bound of $n-c$ registers for long-lived $(k+c)$-adaptive renaming for any integer constant $c \geq 0$.

Our lower bound proofs use covering techniques first introduced by Burns and Lynch [11]. The main challenge is to exploit the semantics of the renaming object to force the processes to write to a large number of registers. In the lower bound for the one-shot case, we first build an execution in which some processes are poised to write to (*cover*) a set of registers. Then we argue that if enough new processes take steps after this, at least one of them must become poised to write to a register not already covered, since, otherwise, the covering processes can obliterate all the traces of the new processes, causing some `getName()` to return an incorrect result. For the lower bound for the long-lived case, we exploit that fact that processes can perform `getName()` and `relName()` repeatedly to build a long execution, where in each inductive step either another register is covered or an available name is used up without being detected by other processes.

2 Preliminaries

This section describes our model of computation and the notation, vocabulary and general techniques used in this paper. Previous work by many researchers (for example [8, 11, 18, 19, 23, 27]) has collectively developed similar tools that serve to make our description of results and presentation of proofs precise, concise and clear. Much of the terminology presented in this section is borrowed or adapted from this previous research.

Our computational model is an asynchronous shared memory system consisting of n processes $\mathcal{P}=\{p_1,\ldots,p_n\}$ and m shared registers $\mathcal{R}=\{R_1,\ldots,R_m\}$. Each process executes code that can access its own independent random number generator and its own private registers as well as the shared registers. Each shared register supports two operations, read and write. Each such operation happens atomically in memory. Processes can only communicate via those operations on shared registers. The algorithm is *deterministic* if each process' code is deterministic; that is, no process' code contains any random choice.

Informally, an execution arises one step at a time, where a step consists of some process, chosen arbitrarily, executing the next instruction in its code. This instruction could be a shared memory access, or a local memory access, or a local operation including, possibly, a random choice. Notice, however, that after a process takes a shared memory step, the outcome of all its subsequent local operations and random choices up to (but not including) its next shared memory operation is independent of any intervening operations by other processes. Therefore, there is no loss of generality in assuming that a step by a process consists of a single shared memory access (or, initially, its method-call invocation) followed by all its subsequent local operations and random choices, up to the point where that process is poised to execute its next shared memory operation.

A *configuration* C is a tuple $(s_1,\ldots,s_n,v_1,\ldots,v_m)$, denoting that process p_i, $1 \leq i \leq n$, is in state s_i and register r_j, $1 \leq j \leq m$, has value v_j. Furthermore the state s_i of p_i is one in which p_i's next operation is either a shared memory operation or an invocation of a method-call (`getName()` or `relName()`). Configurations will be denoted by capital letters. The initial configuration, where each process' next step is to invoke a method-call, is denoted C^*.

Given a configuration, C, a *step from* C is a pair of the form (p,τ) where p is a process identifier, and τ is a sequence of outcomes that arise from the sequence of all random choices that p makes after completing its pending shared mem-

ory operation starting from configuration C up to the point where p is poised to do its next shared memory operation. An execution is an alternating sequence of configurations and steps starting and ending with a configuration, and defined inductively as follows. The 0-step or empty execution starting at C is (C). A k-step execution, $k \geq 1$, starting at C_0 is a sequence $(C_0,(q_1,\tau_1),C_1,(q_2,\tau_2),\ldots,(q_k,\tau_k),C_k)$ where

- $(C_0,(q_1,\tau_1),C_1,(q_2,\tau_2),\ldots,(q_{k-1},\tau_{k-1}),C_{k-1})$ is a $k-1$ step execution starting at C_0, and
- (q_k,τ_k) is a step from C_{k-1} and C_k is the configuration resulting from that step.

An *execution* is a k-step execution for any integer $k \geq 0$. A subsequence, $\sigma = ((q_1,\tau_1),\ldots,(q_k,\tau_k))$, consisting of the steps from an execution starting at C is called a *schedule starting at* C. If σ is a schedule starting at C, then the execution starting at C arising from σ is abbreviated $E = (C;\sigma)$ and $\sigma(C)$ denotes the final configuration of E. If an algorithm is deterministic, then the second component of every step of every execution of the algorithm is empty because there are no random choices. So in this case a schedule is simplified to just a sequence of process identifiers.

A configuration, C, is *reachable* if there exists a finite schedule, σ, such that $\sigma(C^*) = C$. Let σ and π be two finite schedules such that σ starts at configuration C and π starts at $\sigma(C)$. Then $\sigma\pi$ denotes the concatenation of σ and π, and is a schedule starting at C. Let $P \subseteq \mathcal{P}$ be a set of processes, and σ a schedule. We say σ is *P-only* if all the identifiers of processes that appear in σ are in P. If the set P contains only one process, p, then we say σ is *p-only*. We denote the set of processes that appear in schedule σ by $\text{procs}(\sigma)$.

A deterministic implementation of a method is *wait-free* if, for any reachable configuration C and any process p, p completes its method call in a finite number of its own steps, regardless of the steps taken by other processes. An implementation of a method is *non-deterministic solo-terminating* if, for any reachable configuration C and any process p, there exists a finite p-only schedule, σ, starting from C such that p has finished its method call in configuration $\sigma(C)$[18]. Non-deterministic solo-termination for deterministic implementations is called *obstruction-free*.

We say process p *covers* register r in a configuration C, if p writes to r in its next step. A set of processes P covers a set of registers R if for every register $r \in R$ there is a process $p \in P$ such that p covers r. If $|P| = |R|$, then we say P *exactly* covers R. Consider a process set P that exactly covers the register set R in configuration C. Let π_P be any permutation which includes exactly one step by each process in P. Then the execution $(C;\pi_P)$ is called a *block-write* by P to R. Two configurations $C = (s_1,\ldots,s_n,v_1,\ldots,v_m)$ and $C' = (s'_1,\ldots,s'_n,v'_1,\ldots,v'_m)$ are *indistinguishable* to process p_i if $s_i = s'_i$ and $v_j = v'_j$ for $1 \leq j \leq n$. Let P be a set of processes, and σ any P-only schedule starting at configuration C. If for every process $p \in P$, C and C' are indistinguishable to p, then σ is also a schedule starting at C' and $\sigma(C)$ and $\sigma(C')$ are indistinguishable to p.

A process p *participates in configuration* C if in C, p has started a `getName()` operation and has not completed the following `relName()`. A process is called *idle in configuration* C if it does not participate in C. A configuration C is called *quiescent* if, $\forall p \in \mathcal{P}$, p is idle in C. We say process p *owns name* x *in configuration* C if in C, p has completed a `getName()` operation that returned name x and p has not started `relName()`. Let $C_0,\ldots,C_e$ be a sequence of configurations arising from execution E. The number of *participants in* E is the maximum over all i, $0 \leq i \leq e$, of the number of participants in C_i. Given these definitions, the definition of a renaming object can be stated more precisely as follows. Let $f : \{1,\ldots,n\} \to \mathbb{N}$ be a non-decreasing function satisfying $f(1) \leq n-1$. Both long-lived and one-shot b-bounded f-adaptive renaming support the operation `getName()`. Operation `getName()` by process p returns a name x to p. The long-lived version also supports the operation `relName()`, which releases the name x. Both versions must satisfy 1) there is no reachable configuration in which two processes own the same name, 2) if the number of participants during p's `getName()`, k, is at most b then, $x \in \{1,\ldots,f(k)\}$. Observe that, properties 1) and 2) imply that $f(k) \geq k$ for all $k \in \{1,\ldots,n\}$.

3 A Space Lower Bound for Long-Lived Loose Renaming Objects

For any non-decreasing function f satisfying $f(1) \leq n-1$, let d be the largest integer such that $f(d) \leq n-1$. We prove that at least $d+1$ registers are required for non-deterministic solo-terminating long-lived f-adaptive renaming in our system. The proof relies on two lemmas. Lemma 3.1 says that there is no reachable configuration C in which $n-d$ processes own names in the range $\{1,\ldots,n-1\}$ while all of the other d processes are idle and unaware of any of the processes with names. The intuition for this proof is simple: if such a reachable configuration C exists, then there is a configuration reachable from C in which $(n-d)+d = n$ processes all own names in the range $\{1,\ldots,n-1\}$. Lemma 3.2 provides the core of the lower bound argument and the intuition is as follows. Let C be any reachable

configuration in which fewer than $n-d$ processes own names in the range $\{1,\ldots,n-1\}$ while all of the other $d+1$ processes are idle and unaware of the processes with names. Then there is a reachable configuration from C in which either $d+1$ distinct registers are written, or one more name is owned, and the unnamed processes are again idle and still unaware of the processes with names. Since the initial configuration has no processes with names, and all processes are idle, we can apply Lemma 3.2 repeatedly until either we have exactly covered $d+1$ registers or we reach a configuration in which $n-d-1$ processes own names in the range $\{1,\ldots,n-1\}$. Since, according to Lemma 3.1, we cannot get beyond an $(n-d-1)$-invisibly-named configuration, we must eventually exactly cover $d+1$ registers, completing the proof. We will see, in the formal proof, that the result applies even when the renaming implementation is $(d+1)$-bounded.

The definitions and lemmas that follow refer to any non-deterministic solo-terminating implementation from shared registers of a long-lived f-adaptive renaming object. For a configuration C and a set of processes Q, we say Q is *invisible* in C, if there is a reachable quiescent configuration D such that C and D are indistinguishable to all processes in $\overline{Q}$. If the set Q contains only one process q, then we say process q is invisible. Configuration C is called *ℓ-invisibly-named*, if there is a set Q of ℓ processes, such that in C every process in Q owns a name in $\{1,\ldots,n-1\}$ and Q is invisible.

Lemma 3.1. *For the largest integer d satisfying $f(d) \leq n-1$, there is no reachable $(n-\dot{)}$-invisibly-named configuration.*

Proof. By way of contradiction, suppose that there exists a set Q of $n-d$ processes such that in configuration C, all processes in Q are invisible and own names in the range $\{1,\ldots,n-1\}$. Since Q is invisible in C, there is a reachable quiescent configuration D such that D and C are indistinguishable to $\overline{Q}$. Let σ be a $\overline{Q}$-only schedule such that in execution $(D;\sigma)$, all processes in $\overline{Q}$ perform a complete `getName()`. Because $|\overline{Q}| = d$ all processes in $\overline{Q}$ get names in the range $\{1,\ldots,f(d)\} \subseteq \{1,\ldots,n-1\}$. Since C and D are indistinguishable to $\overline{Q}$, all processes in $\overline{Q}$ perform a complete `getName()` during $(C;\sigma)$ and get names in the range $\{1,\ldots,f()\}$ as well. Therefore in configuration $\sigma(C)$ all processes in $Q \cup \overline{Q}$ have names in the range $\{1,\ldots,n-1\}$. However $|Q \cup \overline{Q}| = n$. This is a contradiction because this implies that the acquired names are not distinct. □

The intuition for Lemma 3.2 is as follows. Recall that in an ℓ-invisibly-named configuration, ℓ processes have names, the $n-\ell$ others are idle and unaware of the presence of the invisibly-named processes, and no register is covered. Starting from this configuration we select one process at a time from the set of idle processes and let it execute until either it covers a register not already covered, or it gets a name without covering a new register. We continue this construction as long as the selected process covers a new register. If we reach $d+1$ processes covering distinct registers we are done. Otherwise, we reached a configuration in which one more process holds a name. Furthermore, we can obliterate the trace of this process with the appropriate block write, and then let all other non-idle processes complete their `getName()` methods and the following `relName()`. This takes us to an $(\ell+1)$-invisibly-named configuration.

Lemma 3.2. *Let d be the largest integer such that $f(d) \leq n-1$. For any $0 \leq \ell \leq n-d-1$ and any reachable ℓ-invisibly-named configuration C, there exists a schedule σ, where $|\text{procs}(\sigma)| \leq d+1$, and either*

- *in configuration $\sigma(C)$ at least $d+1$ distinct registers are exactly covered; or*
- *configuration $\sigma(C)$ is $(\ell+1)$-invisibly-named.*

Proof. Let C be an ℓ-invisibly-named configuration, and let Q be the set of ℓ processes that are invisible in C. Let D be a quiescent configuration that is indistinguishable from C for all processes in $\overline{Q}$. First, we inductively construct a sequence of schedules $\delta_0, \delta_1, \ldots$ until we have constructed δ_{last} such that in $\delta_{last}(C)$ either

a) $d+1$ registers are exactly covered, or,

b) $(\ell+1)$ processes own names in $\{1,\ldots,n-1\}$.

We maintain the invariant that for each $i \in \{0,\ldots,last\}$ in configuration $\delta_i(C)$, a set P_i of i processes exactly covers a set L_i of i distinct registers, $P_i \cap Q = \emptyset$, and δ_i is P_i-only. Let δ_0 be the empty schedule. Then in configuration $\delta_0(C) = C$, no register is covered, so the invariant is true for $P_0 = L_0 = \emptyset$.

Now consider $i \geq 0$. If a) or b) holds for δ_i, we let $last = i$ and are done. Otherwise, since in $\delta_i(C)$ a set L_i of i distinct registers is covered, we have $i \leq d$. We construct δ_{i+1} as follows. Select $p \in \overline{P_i \cup Q}$. Let γ be the shortest p-only schedule such that either

1) p does a complete `getName()` in execution $(\delta_i(C);\gamma)$, or
2) in configuration $\gamma(\delta_i(C))$, p covers a register $r \notin L_i$.

Let δ_{i+1} be $\delta_i\gamma$. First assume case 1) happens. By construction the process that performs δ_1 does not write to any register. If $i = 0$ and p does a complete `getName()` in execution $(\delta_0(C);\gamma)$, then $last = 1$ and we are done. For any $i \geq 1$, because Q is invisible to p, in $(\delta_i(C);\gamma)$ p becomes aware of at most the $i-1 \leq d$ other processes in P_i. Since $f(d) \leq n-1$, p gets a name in $\{1,\ldots,n-1\}$, and thus in configuration $\delta_{i+1}(C)$ all processes in $Q \cup \{p\}$ own names in $\{1,\ldots,n-1\}$ and $|Q \cup \{p\}| = \ell+1$. So condition b) is achieved, the construction stops and $\delta_{last} = \delta_{i+1}$.

Now suppose case 2) happens. If $i+1 = d+1$, then condition a) is achieved, the construction stops and $\delta_{last} = \delta_{i+1}$. Otherwise, the invariant remains satisfied for $L_{i+1} = L_i \cup \{r\}$ and $P_{i+1} = P_i \cup \{p\}$. Clearly, after at most $d+1$ steps either a) or b) is achieved.

Now, using schedule δ_{last} we construct schedule σ. If $\delta_{last}(C)$ satisfies a), let $\sigma = \delta_{last}$ and the lemma holds. Hence, suppose that $\delta_{last}(C)$ satisfies b). Let α be the P_{last-1}-only schedule such that in execution $(\delta_{last}\pi_{P_{last-1}}(C);\alpha)$ every process $q \in P_{last-1}$ completes its pending `getName()` operation and performs a complete `relName()`. During execution $(C;\delta_{last})$ only registers in L_{last-1} were written and in configuration $\delta_{last}(C)$, P_{last-1} exactly covers these registers. Because $\delta_{last} = \delta_{last-1}\gamma$ for some p-only postfix γ of δ_{last}, after a block write by P_{last-1}, configurations $\delta_{last}\pi_{P_{last-1}}(C)$ and $\delta_{last-1}\pi_{P_{last-1}}(C)$ are indistinguishable to $\overline{Q \cup \{p\}}$. Since C and D are indistinguishable to $\overline{Q}$, configurations $\delta_{last-1}\pi_{P_{last-1}}(C)$ and $\delta_{last-1}\pi_{P_{last-1}}(D)$ are also indistinguishable to $\overline{Q}$. So, configurations $\delta_{last}\pi_{P_{last-1}}(C)$ and $\delta_{last-1}\pi_{P_{last-1}}(D)$ are indistinguishable to $\overline{Q \cup \{p\}}$. Hence, configurations $\delta_{last}\pi_{P_{last-1}}\alpha(C)$ and $\delta_{last-1}\pi_{P_{last-1}}\alpha(D)$ are indistinguishable to $\overline{(Q \cup \{p\})}$. Since $\delta_{last-1}\pi_{P_{last-1}}\alpha(D)$ is quiescent, configuration $\delta_{last}\pi_{P_{last-1}}\alpha(C)$ is an $(\ell+1)$-invisibly-named configuration. Therefore, the lemma holds for $\sigma = \delta_{last}\pi_{P_{last-1}}\alpha$. □

Theorem 3.3. *Let d be the largest integer such that $f(d) \leq n-1$. Any non-deterministic solo-terminating implementation of a long-lived d-bounded f-adaptive renaming object requires at least $d+1$ registers.*

Proof. Note that C^* is a reachable 0-invisibly-named configuration. We iteratively construct a sequence of schedules $\sigma_0, \sigma_1, \ldots, \sigma_{last}$ as follows: If $0 \leq i \leq n-d$ and C_i is a reachable i-invisibly-named configuration, we apply Lemma 3.2 to obtain a schedule σ_i, $|\text{procs}(\sigma_i)| \leq d+1$, such that $C_{i+1} = \sigma_i(C_i)$ is either an $(i+1)$-invisibly-named configuration, or in C_{i+1} at least $d+1$ distinct registers are covered. In the latter case we let $last = i+1$ and finish the iterative construction. By Lemma 3.1, there is no $(n-d)$-invisibly-named configuration. Hence if the iterative construction reaches a $(n-d-1)$-invisibly-named configuration, by Lemma 3.2, there is a reachable configuration, in which $d+1$ registers are covered. □

Corollary 3.1. *Let $c \in \{1,\ldots,n-1\}$ and $b = n-c$. Any non-deterministic solo-terminating implementation of a long-lived b-bounded $(k+c)$-adaptive renaming object requires at least b registers.*

4 A Space Lower Bound for One-shot Additive Loose Renaming

In one-shot renaming, each process is constrained to call `getName()` at most once (and does not invoke `relName()`), which imposes a severe restriction on the techniques available for proving lower bounds. In particular, constructions that rely on processes repeatedly getting and releasing names cannot be used for one-shot lower bounds. We observed however, that a straightforward linear lower bound for tight renaming actually applies even for one-shot adaptive renaming. Thus, we are motivated to study one-shot renaming objects with looseness constrained by a constant, specifically k-renaming and $(k+c)$-renaming. We refer to one-shot $(k+c)$-renaming object as an additive loose renaming object, where k is the number of participants and $c \geq 0$ is an integer constant. For the case $c = 0$, it is called an adaptive tight renaming object.

Our lower bound proof has a recursive structure and it relies on a generalization of additive loose renaming as follows. For any set $T \subset \{1,\ldots,k+c\}$ where $|T| \leq c$, a $[(k+c)\backslash T]$-renaming object for k processes requires that each participating process returns a unique name from the range $\{1,\ldots,k+c\}\backslash T$.

Lemma 4.1. *Any implementation of $[(k+c)\backslash T]$-renaming uses at least as many registers as an implementation of $[(k+c-|T|)\backslash\emptyset]$-renaming.*

Proof. Let A be a $[(k+c)\backslash T]$-renaming algorithm. Then we construct $[(k+c-|T|)\backslash\emptyset]$-renaming algorithm A' from A without any additional registers as follows. If A returns name x, then A' returns $x-|\{t \in T \mid t \leq x\}|$. Since A returns distinct names in the range $\{1,\ldots,k+c\}\backslash T$, obviously A' returns distinct names in the range $\{1,\ldots,k+c-|T|\}$. □

A process is called *startable in configuration C* if in C, it has not started a `getName()`. Since in one-shot renaming, there is no `relName()` operation, in our proofs in this section instead of using quiescent configurations we are interested in configurations in which each process either has completed its `getName()` operation or it has not started a `getName()`. We call such configurations, *quiet* configurations.

Lemma 4.2. *Let D be a reachable quiet configuration with $n' \geq c+2$ startable processes. For every startable process p, let σ_p denote a p-only schedule such that p performs a complete* `getName()` *in execution $(D;\sigma_p)$. Let Q be any subset of startable processes of size $c+1$, then there exists a process $q \in Q$ such that q writes to a register during $(D;\sigma_q)$.*

Proof. Let X be the set of processes that own names in configuration D. Then processes in X own names in range $\{1,\ldots,|X|+c\}$. Let $Q=\{q_1,\ldots,q_{c+1}\}$. By way of contradiction assume that there is no process $q \in Q$ such that q writes to a register during $(D;\sigma_q)$. Then for all i, $1 \leq i \leq c+1$, configurations $\sigma_{q_1}\ldots\sigma_{q_{c+1}}(D)$ and $\sigma_{q_i}(D)$ are indistinguishable to q_i. Let q' be a startable process not in Q. Hence, $\sigma_{q_1}\ldots\sigma_{q_{c+1}}\sigma_{q'}(D)$ and $\sigma_{q'}(D)$ are indistinguishable to q'. Therefore, all processes in Q plus q' return names from $\{1,\ldots,|X|+1+c\}$ in execution $(D;\sigma_{q_1}\ldots\sigma_{q_{c+1}}\sigma_{q'})$. This is a contradiction because $|X|+c+2$ processes receive names from a set of size $|X|+c+1$ implying that they cannot be assigned distinct names. □

Lemma 4.3. *Let D be a reachable configuration in which:*

- *a set Q of at least $c+1$ processes covers a set of $\ell \geq 1$ registers,*
- *there exists a set $Q' \subseteq Q$ of size $c+1$ such that no process in Q' has written to a register and,*
- *there is a set of $c+1$ startable processes P, disjoint from Q.*

Then, there is a P-only schedule σ_P such that at least $\ell+1$ registers are covered in $\sigma_P(D)$.

Proof. Let X be the set of processes that own names in configuration D. Since no process in Q' has written to a register, processes in X own names in range $\{1,\ldots,|X|+|Q|-|Q'|+c\}$. Let L be the set of registers covered by Q and $Q'' \subseteq Q$ be a set of processes that exactly covers L. Let $\widehat{\sigma_P}$ be a P-only schedule such that in execution $(D;\widehat{\sigma_P})$ all processes in P complete their `getName()` operations. Then all processes in P return names from $\{1,\ldots,|X|+|Q|-|Q'|+|P|+c\} = \{1,\ldots,|X|+|Q|+c\}$ in execution $(D;\widehat{\sigma_P})$. Suppose that in execution $(D;\widehat{\sigma_P})$, there is a process in P that writes to a register not in L. Then let σ_P be the shortest prefix of $\widehat{\sigma_P}$ such that a register $r \notin L$ is covered by a process in P. Hence in configuration $\sigma_P(D)$, L is covered by Q and r is covered by P. Thus we are done. Therefore, assume that in execution $(D;\widehat{\sigma_P})$ all processes in P write only to L. Let $\pi_{Q''}$ be a block-write to L by Q''. Let σ_Q be a Q-only schedule such that in execution $(\sigma_P\pi_{Q''}(D);\sigma_Q)$ all processes in Q complete their `getName()` operations. Since configurations $\sigma_P\pi_{Q''}(D)$ and $\pi_{Q''}(D)$ are indistinguishable to all processes in Q, processes in Q return names from $\{1,\ldots,|X|+|Q|+c\}$ in execution $(\sigma_P\pi_{Q''}(D);\sigma_Q)$. This is a contradiction because $|X|+|Q|+|P| = |X|+|Q|+c+1$ processes receive names from a set of size $|X|+|Q|+c$ implying that they cannot be assigned distinct names. □

Lemma 4.4. *Let D be a reachable configuration in which:*

- *a set P of processes exactly covers a set of $\ell \geq 1$ registers,*
- *there exists a process $q \in P$ such that q has not written to any register and,*
- *there exists a set of $n' \geq c$ startable processes.*

Then, there exists a configuration reachable from D, in which at least $\ell + \lfloor \frac{n'-c}{c+1} \rfloor$ registers are covered.

Proof. Let $\mathcal{P}'$ be the set of all startable processes in D and $Q \subseteq \mathcal{P}'$ be a set of c processes. Then in configuration D, processes in $Q \cup P$ cover a set of ℓ registers where $|Q \cup P| \geq c+1$ and, no process in set $Q \cup \{q\}$ has written to a register. Hence using startable processes in $\mathcal{P}' \setminus Q$, we can inductively apply Lemma 4.3, until we reach a configuration in which $\ell + \lfloor \frac{n'-c}{c+1} \rfloor$ registers are covered. □

Lemma 4.5. *Let A be an non-deterministic solo-terminating implementation of one-shot adaptive tight renaming. Let D be any reachable quiet configuration with $n' \geq 2$ startable processes. Then there is an execution of A, starting from D that requires at least n' registers.*

Proof. Let p be a startable process and $\widehat{\sigma_p}$ be a p-only schedule such that in $(D; \widehat{\sigma_p})$, p completes its `getName()`. Then by Lemma 4.2, p writes to a register. Let σ_p be the shortest prefix of $\widehat{\sigma_p}$ such that in $(D; \sigma_p)$, p covers a register. Then by Lemma 4.4, there exists a configuration reachable from $\sigma_p(D)$, in which at least $1 + n' - 1 = n'$ registers are covered. □

In Lemma 4.6, we show at least $\lfloor \frac{2(n'-c)}{c+2} \rfloor$ registers are required for a non-deterministic solo-terminating implementation of one-shot $(k+c)$-adaptive renaming starting from a quiet configuration with $n' \geq 2c+2$ startable processes. The intuition for this lemma is as follows. We prove the lemma by induction on c. Starting from any quiet configuration, first we select a set Q of $c+1$ startable processes such that one of them writes to a register in a solo-run and we stop it immediately before it writes. Then we choose a process p not in Q and run it until it covers a new register. If we succeed, we select another startable process not in Q. We might not succeed for two reasons. First, we are out of startable processes in which case we are done. Second, process p completes its `getName()` and only writes to the set of covered registers. Then Lemma 4.4 provides a lower bound. Furthermore, starting from this configuration, if the set of covering processes perform a block-write and cover p's trace, then no other process distinguishes this execution from one in which p has not run at all. Therefore by Lemma 4.1, our problem reduces to one-shot $(k+c-1)$-adaptive renaming. Hence we can invoke the induction hypothesis and conclude a second lower bound. Our final lower bound is the maximum of these two lower bounds.

Lemma 4.6. *Let A be a non-deterministic solo-terminating implementation of one-shot $(k+c)$-adaptive renaming. Let D be any reachable quiet configuration with $n' \geq 2c+2$ startable processes. Then there is an execution of A, starting from D that requires at least $\lfloor \frac{2(n'-c)}{c+2} \rfloor$ registers.*

Proof. Let $\mathcal{P}'$ be the set of startable processes in D. We prove the lemma by induction on c. For the base case $c = 0$, by Lemma 4.5, the hypothesis is true. Suppose that the induction hypothesis is true for $c-1 \geq 0$. Since, $|\mathcal{P}'| > c+1$, by Lemma 4.2 there is a process $q \in \mathcal{P}'$ that writes to a register in a solo-execution starting from D. Let σ_q be the shortest q-only schedule such that in configuration $\sigma_q(D)$, there is a register r covered by q. Let $Q \subseteq (\mathcal{P}' \setminus \{q\})$ be a set of c processes.

First, we inductively construct a sequence of schedules $\delta_1, \delta_2, \ldots$ until we have constructed δ_ℓ such that in $\delta_\ell(D)$ either

a) $|\mathcal{P}' \setminus Q|$ registers are covered or,

b) a process q' in $\mathcal{P}' \setminus Q$ has completed its `getName()` and has written only to registers covered by other processes.

We maintain the invariant that for each $i \in \{1, \ldots, \ell\}$ in configuration $\delta_i(D)$, a set $P_i \subseteq (\mathcal{P}' \setminus Q)$ of i processes covers a set L_i of i distinct registers and δ_i is P_i-only.

Let δ_1 be σ_q. Then in configuration $\delta_1(D)$, one register is covered, so the invariant is true for $P_1 = \{q\}$ and $L_1 = \{r\}$.

Now consider $i \geq 1$. If a) or b) holds for δ_i, we let $\ell = i$ and the construction stops. Furtheremore in case b), let q' be the process that completes its `getName()`.

Otherwise, since in $\delta_i(D)$ a set L_i of i distinct registers is covered, we construct δ_{i+1} as follows. Select $p \in \mathcal{P}' \setminus (P_i \cup Q)$. Let γ be the shortest p-only schedule such that either

1) p does a complete `getName()` in execution $(\delta_i(D); \gamma)$ and only writes to L_i, or

2) in configuration $\gamma(\delta_i(D))$, p covers a register $r' \notin L_i$.

Let δ_{i+1} be $\delta_i \gamma$. First assume case 1) happens. Then condition b) is achieved, the construction stops and we let $\delta_\ell = \delta_{i+1}$ and $q' = p$. Now suppose case 2) happens. If $i+1 = |\mathcal{P}' \setminus Q|$, then condition a) is achieved, the construction stops and $\delta_\ell = \delta_{i+1}$. Otherwise, the invariant remains satisfied for $L_{i+1} = L_i \cup \{r'\}$ and $P_{i+1} = P_i \cup \{p\}$. Clearly, after at most $|\mathcal{P}' \setminus Q|$ steps either a) or b) is achieved.

In case a), in configuration $\delta_\ell(D)$, $\ell = n' - c$ registers are covered so the lemma holds in this case. Now suppose case b) happens. In configuration $\delta_{\ell-1}(D)$, a set of $\ell-1$ registers (i.e $L_{\ell-1}$) are covered exactly by a set of processes $P_{\ell-1}$ and process q in $P_{\ell-1}$ has not written to any registers. Furthermore in configuration $\delta_{\ell-1}(D)$, all processes in $\mathcal{P}'\backslash P_{\ell-1} \supseteq Q$ are startable. Therefore, by Lemma 4.4, there is a configuration reachable from $\delta_{\ell-1}(D)$ in which at least $\ell - 1 + \lfloor \frac{n'-\ell+1-c}{c+1} \rfloor$ registers are covered. Let $\pi_{P_{\ell-1}}$ be a block-write by $P_{\ell-1}$. Let $\sigma_{P_{\ell-1}}$ be a $P_{\ell-1}$-only schedule such that in execution $(\delta_\ell \pi_{P_{\ell-1}}(D); \sigma_{P_{\ell-1}})$ all processes in $P_{\ell-1}$ complete their `getName()`. Let x be the name taken by q' in execution $(D; \delta_\ell)$. Let X be the set of processes that own names in configuration D. Since in configuration $\delta_\ell(D)$, only processes in $X \cup P_\ell$ have invoked a `getName()`, $x \in \{1, \ldots, |X| + \ell + c\}$. Note that $\delta_\ell \pi_{P_{\ell-1}} \sigma_{P_{\ell-1}}(D)$ and $\delta_{\ell-1} \pi_{P_{\ell-1}} \sigma_{P_{\ell-1}}(D)$ are indistinguishable to all processes except q'. Hence in any $(\mathcal{P}'\backslash P_\ell)$-only execution starting from $\delta_{\ell-1} \pi_{P_{\ell-1}} \sigma_{P_{\ell-1}}(D)$, names returned by processes in $\mathcal{P}'\backslash P_\ell$ are in $\{1, \ldots, k+c\}\backslash\{x\}$ where $k \geq |X| + \ell$ and therefore $\{x\} \subset \{1, \ldots, k+c\}$. Thus by Lemma 4.1, starting at $\delta_{\ell-1} \pi_{P_{\ell-1}} \sigma_{P_{\ell-1}}(D)$, algorithm A requires as many registers as a $(k+c-1)$-renaming algorithm starting at a quiet configuration with $|\mathcal{P}' - P_\ell| = n' - \ell$ startable processes. Therefore, by the induction hypothesis there is a reachable configuration from $\delta_{\ell-1} \pi_{P_{\ell-1}} \sigma_{P_{\ell-1}}(D)$, in which at least $\lfloor \frac{2(n'-\ell-c+1)}{c+1} \rfloor$ registers are covered.

The minimum of $\ell - 1 + \lfloor \frac{n'-\ell+1-c}{c+1} \rfloor$ and $\lfloor \frac{2(n'-\ell-c+1)}{c+1} \rfloor$, is maximized when $\ell - 1 + \lfloor \frac{n'-\ell+1-c}{c+1} \rfloor = \lfloor \frac{2(n'-\ell-c+1)}{c+1} \rfloor$. Hence, $\ell = \lfloor \frac{n'-c}{c+2} \rfloor - 1$. Therefore the algorithm requires at least $\lfloor \frac{2(n'-c)}{c+2} \rfloor$ registers. □

Theorem 4.7. *Any non-deterministic solo-terminating implementation of one-shot $(k+c)$-adaptive renaming requires at least $\lfloor \frac{2(n-c)}{c+2} \rfloor$ registers.*

Proof. The initial configuration is a quiet configuration with n startable processes. Hence, by Lemma 4.6, there is an execution, starting from the initial configuration that requires at least $\lfloor \frac{2(n-c)}{c+2} \rfloor$ registers. □

Observe that by setting $c = 0$, it follows from Theorem 4.7 that any non-deterministic solo-terminating implementation of one-shot adaptive tight renaming requires n registers. Since the number of startable processes is initially n, next corollary also follows from Lemma 4.6.

Corollary 4.1. *Any non-deterministic solo-terminating implementation of one-shot adaptive tight renaming requires at least n registers.*

5 Wait-Free One-shot $(b-1)$-Bounded $(k(k+1)/2)$-Adaptive Renaming

In this section we present a wait-free one-shot $(b-1)$-bounded $(k(k+1)/2)$-adaptive renaming algorithm from b registers. Since 0-bounded adaptive renaming is a trivial problem, we assume that $b \geq 2$.

The algorithms in this section employ a set $\mathcal{R} = \{R[0], \ldots, R[b-1]\}$ of shared atomic registers. In our proofs, a *register configuration* is a tuple $(V_0, \ldots, V_{b-1})$, denoting that register $R[i]$, $0 \leq i \leq b-1$, has value V_i. The proofs focus on just the sub-sequence of register configurations produced by an execution. Specifically, given an execution $E = (C_0; \sigma)$, let *write schedule* $\widehat{\sigma}$ be the sub-sequence of σ that produces write steps in $(C_0; \sigma)$. Execution E gives rise to the sequence of *register configurations* $\Gamma_E = C_0, C_1, \ldots$ such that the i-th step of $\widehat{\sigma}$ is a write that changes register configuration C_{i-1} to register configuration C_i. For any scan operation s in E, define index$(s) = i$, if s occurs in E between C_i and C_{i+1} in Γ_E. For any write operation w in E, define index$(w) = i$, if w is the i-th step of $\widehat{\sigma}$. Notice that the view returned by a scan with index i is equal to C_i. A register configuration $C = (V_0, \ldots, V_{b-1})$ is *consistent* if $V_0 = \cdots = V_{b-1}$ in which case we say V_0 is the *content* of C. Let C_i and C_j be register configurations in the sequence $\Gamma_E = C_0, C_1, \ldots$ such that $i \leq j$. *Interval*$[i, j]$ denotes the sub-sequence of steps in execution E that begins at write operation w where index$(w) = i$, and ends immediately after write operation u where index$(u) = j$. We use $\nu_C(R)$ to denote the content of register R in configuration C. A local variable x in these algorithms is denoted by x_p when it is used in the method call invoked by process p.

5.1 $(b-1)$-Bounded $(k(k+1)/2)$-Adaptive Renaming Using Atomic Scan

Fig. 1 presents a wait-free implementation of a one-shot $(b-1)$-bounded $(k(k+1)/2)$-adaptive renaming algorithm assuming an atomic scan operation. In Section 5.2, we show how to remove this assumption.

shared: $\mathcal{R} = R[0,\ldots,b-1]$ is an array of multi-writer multi-reader registers, each register is initialized to $\emptyset$
local: An array $r[0,\ldots,b-1]$; $pos \in \{0,\ldots,b-1\}$ initialized to 0; S initialized to $\{id\}$;

Algorithm 1: getName()

```
1  repeat
2  |  R[pos].write(S)
3  |  r[0,...,b-1] := R.scan()
4  |  S := ⋃_{i=0}^{b-1} r[i] ∪ S
5  |  pos := (pos + 1) mod b
6  until (|S| ≥ b) ∨ (r[0] = r[1] = ··· = r[b-1] = S)
7  if |S| ≤ b-1 then
8  |  return (|S|(|S|-1))/2 + rank(id,S)
9  else
10 |  return b(b-1)/2 + id
11 end
```

Figure 1: $(b-1)$-Bounded $(k(k+1)/2)$-Adaptive Renaming Using Atomic Scan

Each process maintains a set of processes, S, that it knows are participating including itself, and alternately executes write and scan operations. In the write operation, it writes S to the next register after where it last wrote, in cyclic order through the b registers. After each of its scan operations, it updates S to all the processes it sees in the scan together with the processes already in its set. The process stops with an assigned name when either its scan shows exactly its own set, S, in every register, or S has grown to size at least b. If $|S|$ is less than b, its name is based on $|S|$ and its rank in S, where $\text{rank}(id,S) = |\{i \mid (i \in S) \wedge (i \leq id)\}|$. If $|S|$ is b or greater, it returns a safe but large name.

Correctness and space complexity.

Since it is clear that the algorithm in Fig. 1 uses b registers, the space complexity will follow immediately after confirming that it is a correct adaptive renaming algorithm. The correctness of this algorithm relies on the fact that if any two processes return names based on a set of size $s < b$, then they have the same set. The main component of the proof is to establish this fact.

Observation 5.1. *For any write operation with value S by process p, $p \in S$.*

Lemma 5.2. *For any execution E, let C_a be a consistent register configuration with content $\widehat{S}$. For any register configuration D following C_a in E, define $\mathcal{T}_D = \{R \in \mathcal{R} | \widehat{S} \not\subseteq \nu_D(R)\}$. Then there exists a one-to-one function $f_D : \mathcal{T}_D \to \mathcal{P}$ satisfying, $\forall R \in \mathcal{T}_D$, $f_D(R) \in \nu_D(R)$ and $f_D(R)$ performs at least one write in the execution interval between C_a and D.*

Proof. Let $C_a, C_{a+1}, \ldots$ be the sequence of register configurations that arises from E starting at C_a. We prove the lemma by induction on the indices of this sequence. The base case $\ell = a$, is trivially true since set $\mathcal{T}_{C_a}$ is the empty set.

Suppose that the induction hypothesis is true for $\ell - 1 \geq a$. Let the write step between $C_{\ell-1}$ and C_ℓ be the operation, w, by process p, into register $\widehat{R}$ with value V. Let s be the most recent scan operation by p preceding w if it exists.

If $\widehat{S} \subseteq V$, then $\mathcal{T}_{C_\ell} = \mathcal{T}_{C_{\ell-1}} \setminus \{\widehat{R}\}$. Define $f_{C_\ell}(R) = f_{C_{\ell-1}}(R)$, $\forall R \in \mathcal{T}_{C_\ell}$. Since $f_{C_{\ell-1}}$ satisfies the induction hypothesis, and $\mathcal{T}_{C_\ell} \subseteq \mathcal{T}_{C_{\ell-1}}$, f_{C_ℓ} also satisfies the induction hypothesis.

Now consider the case $\widehat{S} \not\subseteq V$. So $\mathcal{T}_{C_\ell} = \mathcal{T}_{C_{\ell-1}} \cup \{\widehat{R}\}$. We first show that s happens before C_a or w is the first write by p. Suppose, for the purpose of contradiction, that $a \leq \text{index}(s) \leq \ell - 1$. We have $\forall R \in \mathcal{R}$, $\widehat{S} \not\subseteq \nu_{C_{\text{index}(s)}}(R)$ since otherwise, by Line 4, $\widehat{S} \subseteq V$. Thus $|\mathcal{T}_{C_{\text{index}(s)}}| = b$. By the induction hypothesis, $f_{C_{\text{index}(s)}}$ selects a distinct process from each register in $\mathcal{T}_{C_{\text{index}(s)}}$, implying, by Line 4, that the size of S_p is at least b. Hence p would have stopped in Line 6 before performing any write operation. Therefore s happens before C_a or w is the first write by p, and

consequently any write by p before w happens before C_a. On the other hand, $\forall R \in \mathcal{T}_{C_{\ell-1}}$, $f_{C_{\ell-1}}(R)$ performs a write during Interval$[a, \ell-1]$ implying p is not in $\{f_{C_{\ell-1}}(R) \mid R \in \mathcal{T}_{C_{\ell-1}}\}$. By Observation 5.1, $p \in V$ and p performs a write after C_a. Therefore by defining $f_{C_\ell}(R) = f_{C_{\ell-1}}(R)$, $\forall R \in (\mathcal{T}_{C_{\ell-1}} \setminus \{\widehat{R}\})$ and $f_{C_\ell}(\widehat{R}) = p$, the induction hypothesis holds for ℓ. □

Lemma 5.3. *For any execution E, let $\widehat{S_p}$ and $\widehat{S_q}$ be the value of S_p and S_q in Line 7 for p and q when they have completed the repeat loop. If $|\widehat{S_p}| = |\widehat{S_q}| < b$ then $\widehat{S_p} = \widehat{S_q}$.*

Proof. Let C_p and C_q be the consistent register configurations that resulted in $\widehat{S_p}$ and $\widehat{S_q}$ respectively and assume, without loss of generality, that C_p precedes C_q in Γ_E. By Line 6, $R[0] = \cdots = R[b-1] = \widehat{S_q}$ in C_q. Thus, either $\forall R \in \mathcal{R}$, $\widehat{S_p} \subseteq \nu_{C_q}(R)$ or $\forall R \in \mathcal{R}$, $\widehat{S_p} \not\subseteq \nu_{C_q}(R)$.

For the first case, by Line 4, $\widehat{S_p} \subseteq \widehat{S_q}$ and since $|\widehat{S_p}| = |\widehat{S_q}|$, $\widehat{S_p} = \widehat{S_q}$. For the latter case, set $\mathcal{T}_{C_q} = \{R \in \mathcal{R} \mid \widehat{S_p} \not\subseteq \nu_{C_q}(R)\}$ has size b. By Lemma 5.2, there is a distinct process in each register in $\mathcal{T}_{C_q}$. So there are at least b distinct processes in $\widehat{S_q}$ contradicting $|\widehat{S_q}| < b$. □

Lemma 5.4. *The names returned by any two distinct processes are distinct.*

Proof. Let $\widehat{S_p}$ and $\widehat{S_q}$ be the values of S_p and S_q in Line 7. Without loss of generality, assume that $|\widehat{S_p}| \leq |\widehat{S_q}|$. If $|\widehat{S_p}| \geq b$ and $|\widehat{S_q}| \geq b$, the names returned by p and q in Line 10 are distinct because $p \neq q$. If $|\widehat{S_p}| < b$ and $|\widehat{S_q}| \geq b$, then, by Line 8, the name returned by p is at most $(b-1)(b-2)/2 + (b-1) = b(b-1)/2$ and, by Line 10, the name returned by q is bigger than $b(b-1)/2$. If $|\widehat{S_p}| < b$ and $|\widehat{S_q}| < b$, both processes return at Line 8. First suppose $\ell = |\widehat{S_p}| < |\widehat{S_q}|$. Then the name returned by p is at most $(\ell+1)(\ell)/2$ and the name returned by q is at least $(\ell+1)(\ell)/2 + 1$. If $|\widehat{S_p}| = |\widehat{S_q}|$, by Lemma 5.3, $\widehat{S_p} = \widehat{S_q}$. Therefore $\text{rank}(p, \widehat{S_p}) \neq \text{rank}(q, \widehat{S_q})$. Thus, in all cases the names returned by p and q are distinct. □

Observation 5.5. *Set $\{p\}$ is written by p before any other write of any set $V \supseteq \{p\}$.*

Lemma 5.6. *Let k be the number of participating processes during process p's* `getName()`. *Then, any name returned by p is in the range $\{1, \ldots, \frac{k(k+1)}{2}\}$ if $k < b$ and in the range $\{1, \ldots, n + \frac{b(b-1)}{2}\}$ if $k \geq b$.*

Proof. By Observation 5.5, $\forall q \in S_p$, q performs at least one write before p returns. Thus, $\forall q \in S_p$, q is a participating process. Hence, $|S_p| \leq k$. If $k < b$, then $|S_p| < b$. Therefore, process p returns in Line 8, and the name is in the range $\{1, \ldots, \frac{k(k+1)}{2}\}$. If $k \geq b$, then p returns either in Line 8 or in Line 10. Therefore the name is in the range $\{1, \ldots, \frac{b(b-1)}{2} + n\}$. □

In summary, Lemmas 5.4 and 5.6 imply that the algorithm in Fig. 1 is an $(b-1)$-bounded $(k(k+1)/2)$-adaptive renaming algorithm that uses b registers assuming the availability of the atomic scan operation.

Step complexity.

We now bound the maximum number of steps (scans and writes) that any process can take during its execution of `getName`. Lemma 5.7 establishes the most important piece of the step complexity of the algorithm in Fig. 1. In this lemma we prove that if there exists a register configuration in which there are at least k registers, each of which contains a set of size at least k, then the number of distinct process names in any subsequent scan is at least k. We call such a register configuration *k-complete*, and any set of such registers is a *k-full-set*. The core idea is that after a k-complete configuration with k-full-set $\mathcal{R}'$, every write with set size less than k to a register in $\mathcal{R}'$ is performed by a distinct writer. It then follows that the union of the sets appearing in $\mathcal{R}'$ always will have size at least k. For the proof, given set of registers $\mathcal{R}' \subseteq \mathcal{R}$ and a register configuration D, we will be interested in those registers in $\mathcal{R}'$ that contain a set smaller than k, and in the processes that wrote these small sets to these registers. Therefore, define $\rho_{\mathcal{R}'}(D,k) = \{R \in \mathcal{R}' \mid |\nu_D(R)| < k\}$. Let *writer*$(j,R)$ denote the process that performs the most recent write to register R preceding register configuration C_j. For any set of registers $\mathcal{R}'$, register configuration C_j and an integer k, define $W_{\mathcal{R}'}(j,k) = \{\text{writer}(j,R) \mid R \in \rho_{\mathcal{R}'}(C_j,k)\}$. Notice that a register configuration D is k-complete if there exists a set $\mathcal{R}'$ of k registers where $\rho_{\mathcal{R}'}(D,k) = \emptyset$. Furthermore, $\mathcal{R}'$ is k-full-set at register configuration D.

Lemma 5.7 uses a proof structure that is more elaborate than, but reminiscent of, that of Lemma 5.2.

Lemma 5.7. *For any execution E, let C_a be a k-complete register configuration where $0 \leq k \leq b-1$ and let $\mathcal{R}'$ be a k-full-set of C_a. For any register configuration C_e following C_a in Γ_E, there exists a one-to-one and onto function $g_{C_e} : \rho_{\mathcal{R}'}(C_e,k) \to W_{\mathcal{R}'}(e,k)$ satisfying, $\forall R \in \rho_{\mathcal{R}'}(C_e,k)$, $g_{C_e}(R) \in \nu_{C_e}(R)$. Furthermore, each process in $W_{\mathcal{R}'}(e,k)$ performs at least one write in Interval$[a,e]$.*

Proof. Let $C_a, C_{a+1}, \ldots$ be the sequence of all register configurations starting at C_a. We prove the lemma by induction on the indices of this sequence. The base case $\ell = a$, is trivially true since set $\rho_{\mathcal{R}'}(C_a,k)$ is the empty set.

Suppose that the induction hypothesis is true for $\ell - 1 \geq a$. Let the write step between $C_{\ell-1}$ and C_ℓ be the operation, w, by process q, into register $\widehat{R}$ with value V. Let s be the most recent scan operation by q preceding w if it exists.

Suppose that $\widehat{R} \notin \mathcal{R}'$. Then $\rho_{\mathcal{R}'}(C_{\ell-1},k) = \rho_{\mathcal{R}'}(C_\ell,k)$ and $W_{\mathcal{R}'}(\ell-1,k) = W_{\mathcal{R}'}(\ell,k)$, so the induction hypothesis holds trivially for ℓ by setting $g_{C_\ell} = g_{C_{\ell-1}}$.

Suppose that $\widehat{R} \in \mathcal{R}'$ and $|V| \geq k$. Then $\rho_{\mathcal{R}'}(C_\ell,k) = \rho_{\mathcal{R}'}(C_{\ell-1},k) \setminus \{\widehat{R}\}$ and $W_{\mathcal{R}'}(\ell,k) = W_{\mathcal{R}'}(\ell-1,k) \setminus \{\text{writer}(\ell-1,\widehat{R})\}$. So the hypothesis holds for ℓ by setting $g_{C_\ell} = g_{C_{\ell-1}}$ for each $R \in \rho_{\mathcal{R}'}(C_\ell,k)$.

Finally, consider the case $\widehat{R} \in \mathcal{R}'$ and $|V| < k$. We first show that s happens before C_a or w is the first write by q. Suppose, for the purpose of contradiction, that s happens after C_a. Then $a \leq \text{index}(s) \leq \ell - 1$. For each $R \in \mathcal{R}'$, we have $|\nu_{C_{\text{index}(s)}}(R)| < k$ since otherwise, by Line 4, $|V| \geq k$. Thus $|\rho_{\mathcal{R}'}(C_{\text{index}(s)},k)| = k$. By the induction hypothesis, $g_{C_{\text{index}(s)}}$ is a bijection, so $|\rho_{\mathcal{R}'}(C_{\text{index}(s)},k)| = |W_{\mathcal{R}'}(\text{index}(s),k)| = k$, and $\forall R \in \rho_{\mathcal{R}'}(C_{\text{index}(s)},k)$, $g_{C_{\text{index}(s)}}(R) \in \nu_{C_{\text{index}(s)}}(R)$. Therefore, by Line 4, the size of S_q, and hence the size of V, is at least k, which is a contradiction.

Therefore s happens before C_a or w is the first write by q, and consequently any write by q before w happens before C_a. On the other hand, by the induction hypothesis, $\forall R \in \rho_{\mathcal{R}'}(C_{\ell-1},k)$, writer$(\ell-1,R)$ performs a write during Interval$[a,\ell-1]$ implying q is not in $W_{\mathcal{R}'}(\ell-1,k)$. We have $\rho_{\mathcal{R}'}(C_\ell,k) = \rho_{\mathcal{R}'}(C_{\ell-1},k) \cup \{\widehat{R}\}$ and $W_{\mathcal{R}'}(\ell,k) = W_{\mathcal{R}'}(\ell-1,k) \setminus \text{writer}(\ell-1,\widehat{R}) \cup \{q\}$, whether or not $\widehat{R}$ is in $\rho_{\mathcal{R}'}(C_{\ell-1},k)$. Furthermore, q performs a write after C_a. Therefore, the induction hypothesis holds for ℓ by defining $g_{C_\ell} = g_{C_{\ell-1}}$ for each $R \in \rho_{\mathcal{R}'}(C_\ell,k) \setminus \{\widehat{R}\}$, and $g_{C_\ell}(\widehat{R}) = q$. □

Lemma 5.8. *Let D be a k-complete register configuration. Then for each process p in $\mathcal{P}$, p's second write after D if it exists, has a value with size at least k.*

Proof. Let w be the second write operation by p after D if it exists. Let s be the most recent scan operation by p preceding w. Since w is the second write by p, the value returned by s is equal to a register configuration D' following D. Let $\mathcal{R}'$ be the k-full-set of D. If $\exists R \in \mathcal{R}$, $|\nu_{D'}(R)| \geq k$, then by Line 4, the size of S_p at s is at least k. Otherwise, all registers, and hence all registers in $\mathcal{R}'$, contain sets of size less than k. Therefore, $|\rho_{\mathcal{R}'}(D',k)| = k$. So, by Lemma 5.7, $|W_{\mathcal{R}'}(\text{index}(D'),k)| = k$. Thus, again by Line 4, the size of S_p at s is at least k. Hence w has a value with size at least k.

□

Lemma 5.9. *Let E be an execution whose first operation is a write by p and contains the next b scans by p. Furthermore, during E, every write by p has value Q and no write has value $Q' \subsetneq Q$. Then p either terminates or the size of S_p increases.*

Proof. By Lines 2-3, each scan operation is preceded by a write operation. Hence E contains b writes by p. Therefore, during E, p writes Q to all b registers. Let $(V_0,\ldots,V_{b-1})$ be the value returned by p's last scan during E. Because E does not contain any write with value $Q' \subsetneq Q$ either $V_0 = \cdots = V_{b-1} = Q$ in which case p terminates or $\exists i$, $0 \leq i \leq b-1$ such that $V_i \not\subseteq Q$. In the latter case by Line 4, the size of S_p increases.

□

Lemma 5.10. *Let D be a k-complete register configuration where $0 \leq k \leq b-1$. Then for each process p in $\mathcal{P}$, p's $(bk+2)$-nd write after D, if it exists, has a value with size at least $k+1$.*

Proof. Let w be the second write by p after D. Suppose that p writes Q at w. By Lemma 5.8, $|Q| \geq k$. If $|Q| \geq k+1$ or p terminates before writing bk more times, we are done. Therefore, suppose that $|Q| = k$ and p performs bk writes after w. Then $|Q \setminus \{p\}| = k-1$. By Lemma 5.8, after D, $\forall q \in Q$, q writes a value with size smaller than k at most once.

Let E be the execution whose first operation is w and contains the next bk scan operations by p. Partition E into disjoint segments, $E = (E_1, \ldots, E_k)$, satisfying $\forall \ell$, $1 \leq \ell \leq k$, the first operation in E_ℓ is a write operation by p and E_ℓ contains the next b scans by p. Notice that E contains exactly bk write operations by p and since w is the first operation of E, p performs at least one more write after E ends. Since there are at most $|Q \setminus \{p\}| = k-1$ writes after w that have a value V satisfying $V \subsetneq Q$, there exists an ℓ, $1 \leq \ell \leq k$ such that all writes during E_ℓ have a value that is not a proper subset of Q. Since p does not terminate during E_ℓ, by Lemma 5.9, the size of S_p after E_ℓ (hence, after E) is at least $k+1$. Hence, p's $(bk+2)$-nd write after D has a value with size at least $k+1$.

□

Lemma 5.11. *For any execution E in which p does not terminate, let O be the set of all scan operations by p during E. Let $Z = \{writer(\text{index}(s),R) \mid s \in O \text{ and } R \in \mathcal{R}\}$. Then $|Z| < b$.*

Proof. For any scan $s \in O$ and any register R, $\text{writer}(\text{index}(s),R) \in \nu_{C_{\text{index}(s)}}(R)$. Therefore, by Line 4, for any $s \in O$, after s, S_p contains $\text{writer}(\text{index}(s),R)$. Since p does not terminate after s, at s, $|S_p| < b$. Hence $|Z| < b$.

□

Lemma 5.12. *Let D be a k-complete register configuration where $0 \leq k \leq b-1$. Then for each process p in $\mathcal{P}$, p makes at most $bk+1+b(\frac{(b-1)(bk+1)}{b-k}+1)$ write operations before it terminates or a $(k+1)$-complete register configuration is achieved.*

Proof. By Lemma 5.11, p's $(bk+2)$-nd write, say w, after D has a value with size at least $k+1$. Let E be the execution whose first operation is w and contains the next $b(\frac{(b-1)(bk+1)}{b-k}+1)$ scan operations by p. Partition E into segments, $E = (E_1, \ldots, E_{\frac{(b-1)(bk+1)}{b-k}+1})$, satisfying $\forall \ell$, $1 \leq \ell \leq \frac{(b-1)(bk+1)}{b-k}+1$, the first operation in E_ℓ is a write operation by p and E_ℓ contains the next b scans by p. Let O be the set of all scan operations by p during E. Let $Z = \{\text{writer}(\text{index}(s),R) \mid s \in O \text{ and } R \in \mathcal{R}\}$. Let U be the set of all write operations by processes in Z during E such that $\forall u \in U$, the value of u has a size smaller than or equal to k. By Lemmas 5.10 and 5.11, $|U| \leq (bk+1)|Z| \leq (b-1)(bk+1)$. Let $U_\ell = \{u \mid u \in U \text{ and } u \text{ happens during } E_\ell\}$.

By the pigeon whole principle, there exists an ℓ such that $|U_\ell| < b-k$. Let s_ℓ be p's last scan during E_ℓ. Since during E_ℓ, p writes a value with size at least $k+1$ to all b registers and the number of writes with value smaller than $k+1$ and scanned by p (i.e $|U_\ell|$), is less than $b-k$, s_ℓ returns a view in which at least $k+1$ registers have size at least $k+1$. Hence, $C_{\text{index}(s_\ell)}$ is $(k+1)$-complete.

□

Lemma 5.13. *No process writes more than $3b^4 \ln b$ times.*

Proof. By Lemma 5.12, a process can write at most $bk+1+b(\frac{(b-1)(bk+1)}{b-k}+1)$ times between a k-complete and a $(k+1)$-complete configuration. The initial configuration is 0-complete and an b-complete configuration cannot exist. Therefore a process can write a most

$$\begin{aligned}
& \sum_{k=0}^{b-1} \left(bk+1+b(\frac{(b-1)(bk+1)}{b-k}+1)\right) \\
= & \sum_{k=0}^{b-1} (bk+1+b) + b(b-1)\sum_{k=0}^{b-1} (\frac{bk+1}{b-k}) \\
= & \sum_{k=0}^{b-1} (bk+1+b) + b(b-1)\sum_{k=1}^{b} (\frac{b^2-bk+1}{k}) \\
< & \frac{b^3}{2} + b - 1 + b^2 + b^2 \sum_{k=1}^{b} (\frac{b^2}{k} - b + \frac{1}{k}) < 3b^4 \ln b
\end{aligned}$$

times before it terminates.

□

5.2 $(b-1)$-Bounded $(k(k+1)/2)$-Adaptive Renaming Using Registers

We replace the atomic scan in Fig. 1 with a new function, `newScan()`, and the `getName()` algorithm also changes accordingly. The revised renaming algorithm is shown in Fig. 2. In the `getName()` algorithm, processes augment the values they write to each register with their ids and sequence numbers in order to guarantee the uniqueness of the value of each write. This prevents the ABA problem. Each register $R \in \mathcal{R}$ stores an ordered triple $(set, id, seqNumber)$.

During a `newScan()` operation by process p, p performs a `collect(`$\mathcal{R}$`)` in Line 18, by reading $R[0]$ through $R[b-1]$ consecutively and returns a *collect*. After each `collect(`$\mathcal{R}$`)`, p updates its set S from this collect. It repeatedly gets a collect until either the size of set S becomes at least b or it obtains two identical consecutive collects and returns this collect. If `newScan()` terminates at Line 24, then the returned collect is equivalent to the returned value of a linearizable implementation of a scan [21, 26, 25, 14, 15]. Hence all the proofs in Section 5.1 hold when `newScan()` terminates at Line 24. So to establish the correctness, we need to prove that when a process p returns in Line 21, in fact more than $b-1$ processes are participating, hence the name returned by p's `getName()` is valid. This is shown in Lemma 5.14. In Lemma 5.13, we showed that the number of writes by each process is bounded. Since the sequence number *seqNumber*, cannot get larger than the number of write operations by each process, the size of each register is also bounded. Therefore, after Lemma 5.14 it will remain to prove that the `getName()` algorithm in Fig. 2 is wait-free. This will be established in Lemmas 5.15 through 5.17, by bounding the number of steps of each `newScan()` operation.

For any read operation o of register R, define writeOp(o) to be the most recent write operation to R preceding o if it exists and $\perp$ otherwise. For any write operation w, let performer(w) denote the process that performs w. For any set of write operations W, let $\mathcal{Z}(W) = \{\text{performer}(w) \mid w \in W\}$.

Lemma 5.14. *Let k be the number of participating processes during process p's* `getName()`. *If a* `newScan()` *operation by p returns in Line 21, then $k \geq b$.*

Proof. Let $\widehat{S_p}$ be the value of S_p when p's `getName()` returns. Since p's `newScan()` operation returns in Line 21, $|\widehat{S_p}| \geq b$. By Observation 5.5, $\forall q \in \widehat{S_p}$, q performs at least one write before p returns. Thus, $\forall q \in \widehat{S_p}$, q has invoked a `getName()` before p returns. Therefore, $k \geq |\widehat{S_p}| \geq b$. □

Lemma 5.15. *Let E be an execution such that any step by process p during E is part of a single* `newScan()` *operation. If E contains at least $3b$ reads by p and does not contain any write operations, then p's* `newScan()` *terminates during E.*

Proof. Since E contains no write operation, every $3b$ reads by p must contain two complete identical collects. Hence, p must terminate due to Line 23.

□

Lemma 5.16. *Let E be an execution such that any step by process p during E is part of a single* `newScan()` *operation, s. Let O be the set of all reads by p during E and $W = \{\text{writeOp}(o) \mid o \in O\} \backslash \{\perp\}$. If $|\mathcal{Z}(W)| \geq b$, then s contains at most $2b$ read operations after E ends.*

Proof. If p performs fewer than $2b$ read operations after E ends, we are done. Let E' be an execution which starts after E ends and contains $2b$ reads by p. Since every $2b$ reads by p must contain a complete collect, after p's complete collect during E', by Line 19, S_p includes all processes in $\mathcal{Z}(W)$. Hence, after p's complete collect during E', $|S_p| \geq |\mathcal{Z}(W)| \geq b$. Therefore, by Line 20, s must terminate. □

Lemma 5.17. *No* `newScan()` *operation contains more than $10b^6 \ln b$ reads.*

Proof. By way of contradiction, let E be an execution in which process p performs a single `newScan()` s, and it contains more than $10b^6 \ln b$ reads. Let E' be a prefix of E that contains $9b^6 \ln b$ reads by p. Partition E' into disjoint segments, $E' = (E_1, \ldots, E_{3b^5 \ln b})$, satisfying $\forall \ell$, $1 \leq \ell \leq 3b^5 \ln b$, E_ℓ contains $3b$ reads by p. Let O be the set of all read operations by p during E' and $W = \{\text{writeOp}(o) \mid o \in O\} \backslash \{\perp\}$.

Suppose there is an ℓ such that E_ℓ contains no write operation in W. This implies that E_ℓ contains no write operation. Therefore by Lemma 5.15, p terminates s during E_ℓ.

shared: $R[0\ldots b-1]$ is an array of multi-writer multi-reader registers, each register is initialized to $(\emptyset,0,0)$;
local: $r[0,\ldots,b-1]$; $pos \in \{0,\ldots,b-1\}$ initialized to 0; *seqNumber* is a non-negative integer initialized to 0; S is initialized to $\{id\}$; *largeSet* is a boolean;

Algorithm 2: getName()

```
1  repeat
2      seqNumber = seqNumber + 1
3      R[pos].write(S, id, seqNumber)
4      (largeSet, r[0,...,b-1]) := R.newScan(S)
5      if largeSet then
6          return b(b-1)/2 + id
7      end
8      S := ⋃_{i=0}^{b-1} r[i].set
9      pos := (pos + 1) mod b
10 until (|S| ≥ b) ∨ (r[0].set = r[1].set = ··· = r[b-1].set = S)
11 if |S| ≤ b-1 then
12     return (|S|(|S|-1))/2 + rank(id, S)
13 else
14     return b(b-1)/2 + id
15 end
```

local: $a[0\ldots b-1]$; $a'[0\ldots b-1]$ each element is initialized to $(\emptyset,0,0)$;

Algorithm 3: newScan(S)

```
16 repeat
17     a := a'
18     a' := collect(R)
19     S := ⋃_{i=0}^{b-1} a'[i].set ∪ S
20     if |S| > b-1 then
21         return (True, a')
22     end
23 until a = a'
24 return (False, a')
```

Figure 2: $(b-1)$-Bounded $(k(k+1)/2)$-Adaptive Renaming Using Registers

Otherwise, each E_ℓ contains at least one write in W. Hence $|W| \geq 3b^5 \ln b$. Since by Lemma 5.13, each process p performs at most $3b^4 \ln b$ writes, $|\mathcal{Z}(W)| \geq b$. Therefore, by Lemma 5.16, s contains at most $2b$ reads after E' ends. Hence E contains at most $9b^6 \ln b + 2b < 10b^6 \ln b$ reads by p. □

Lemma 5.18. *No process performs more than* $31b^{10} \ln^2 b$ *shared steps (read or write).*

Proof. By Lemma 5.13, each process p performs at most $3b^4 \ln b$ writes. Hence, p performs at most $3b^4 \ln b$ `newScan()` operations. By Lemma 5.17, p performs at most $10b^6 \ln b$ reads in each `newScan()` operation. Hence p performs at most $3b^4 \ln b + (3b^4 \ln b)(10b^6 \ln b) \leq 31b^{10} \ln^2 b$ shared steps. □

Theorem 5.19. *For any $b \geq 2$, there is a wait-free one-shot $(b-1)$-bounded $(k(k+1)/2)$-adaptive renaming algorithm implemented from b bounded registers. Additionally, when $k \geq b$, the returned names are in the range $\{1, \ldots, n + \frac{b(b-1)}{2}\}$.*

Setting $b = \lceil \sqrt{n} \rceil + 1$, we have a wait-free one-shot $\lceil \sqrt{n} \rceil$-bounded $(k(k+1)/2)$-adaptive renaming algorithm from $\lceil \sqrt{n} \rceil + 1$ bounded registers. This implies that the algorithm returns names in the range $\{1, \ldots, (k(k+1)/2)\}$ when $k \leq \lceil \sqrt{n} \rceil$, and returns names in the range $\{1, \ldots, n + \frac{\lceil \sqrt{n} \rceil (\lceil \sqrt{n} \rceil + 1)}{2}\}$ when $k \geq \lceil \sqrt{n} \rceil + 1$. Note that when $k \geq \lceil \sqrt{n} \rceil + 1$, $n + \frac{\lceil \sqrt{n} \rceil (\lceil \sqrt{n} \rceil + 1)}{2} \leq k^2 + \frac{k^2}{2}$. Hence, $\forall k \in \{1, \ldots, n\}$, the algorithm returns names in the range $\{1, \ldots, (3k^2)/2\}$.

Corollary 5.1. *There is a wait-free one-shot $(3k^2)/2)$-adaptive renaming algorithm implemented from $\lceil \sqrt{n} \rceil + 1$ bounded registers.*

6 Obstruction-Free $(b-1)$-Bounded k-Adaptive Renaming

Fig. 3 presents pseudo-code for an obstruction-free one-shot $(b-1)$-bounded k-adaptive renaming algorithm from b registers assuming an atomic scan operation. In Theorem 6.7, we show how to remove this assumption by adding an extra register.

Algorithm Description.

A *naming set* is a set of ordered pairs where each pair is a process id and a proposed name with the property that no process id occurs in more than one pair in the set. Let S be a naming set. In our algorithm and the analysis we use the following notation:

- $\text{Procs}(S) = \{x \mid (x,y) \in S\}$,
- $\text{Names}(S) = \{y \mid (x,y) \in S\}$,
- if $(p,n) \in S$, then $\text{name}(p,S)$ is n; otherwise it is undefined.

The algorithm in Fig. 3 employ a set $\mathcal{R} = \{R[1], \ldots, R[b]\}$ of shared atomic registers. Each register R stores an ordered triple $(set, writer, proposal)$ where *set* is a naming set, *writer* is a process id or $\perp$ (initially) and *proposal* is a positive integer less than or equal $b-1$. Each process p maintains a naming set S_p and alternates between write and scan operations until it terminates with a name for itself. Each scan returns a view, which is an atomic snapshot of the content of all registers. Each write by p writes a triple consisting of its set S_p, its *id* p, and its proposed name $name_p$, to some register $R[j]$. Process p uses its last view and its previous value of S_p to determine the new value of S_p, $name_p$ and j.

Function `Update` describes how p constructs S_p in three steps. In the first step (Lines 18-22), p creates a naming set based only on the *writer*s and *proposal*s of each register in its view. If the view contains a *writer* with more than one *proposal*, p chooses one pair arbitrarily. In the second step (Lines 23-31), p augments its naming set with additional pairs for processes that are not *writer*s in its view but occur in the union of all naming sets in its view. The main issue occurs when there is some process that is paired with more than one name from two or more naming sets in different registers. In this case, if there are two such registers with the same *writer* then, p chooses the pair which occurs in the register with bigger index. Otherwise, p picks one pair arbitrarily. Finally (Lines 32-34), p adds any pair (q, n_q) such that q exists in the previous version of S_p and is not yet added. Observe that S_p is a naming set and $p \notin \text{Procs}(S_p)$.

In Line 5 p chooses its proposal for its own name, $name_p$, to be the smallest integer that is not paired with some other process in S_p.

shared: $\mathcal{R} = R[1,\ldots,b]$ is an array of multi-writer multi-reader registers, each register is initialized to $(\emptyset,\perp,1)$
local: $r[1,\ldots,b]$; $pos \in \{1,\ldots,b\}$ initialized to 1; S initialized to $\emptyset$; $proposed \in \mathbb{N}$ initialized to 1.

Algorithm 4: getName()

```
1  repeat
2      R[pos].write(S, id, proposed)
3      r[1,...,b] := R.scan()
4      S := Update(S, r[1,...,b])
5      proposed = min{i ∈ ℕ | i ∉ Names(S)}
6      if ∃i, s.t. (r[i].writer = id) ∧ (r[i] ≠ (S, id, proposed)) then
7          pos := max{i | (r[i].writer = id) ∧ (r[i] ≠ (S, id, proposed))}
8      else if ∃j, s.t. r[j] ≠ (S, id, proposed) then
9          pos := j
10     end
11 until (|S| + 1 ≥ b) ∨ (r[1] = r[2] = ··· = r[b] = (S, id, proposed))
12 if |S| + 1 ≤ b − 1 then
13     return proposed
14 else
15     return b − 1 + id
16 end
```

Algorithm 5: Update()

```
17 S_new = ∅
18 for all w ∈ {r[i].writer | 1 ≤ i ≤ b} \ {id, ⊥} do
19     Let j ∈ {1,...,b} such that r[j].writer = w
20     name_w := r[j].proposal
21     S_new := S_new ∪ {(w, name_w)}
22 end
23 for ∀p ∈ Procs(⋃_{i=1}^{b} r[i].set) \ (Procs(S_new) ∪ {id}) do
24     if ∃i, j, (i < j) ∧ (r[i].writer = r[j].writer) ∧ (p ∈ Procs(r[j].set)) then
25         name_p := name(p, r[j].set)
26     else
27         Let j ∈ {1,...,b} s.t. p ∈ Procs(r[j].set)
28         name_p := name(p, r[j].set)
29     end
30     S_new := S_new ∪ {(p, name_p)}
31 end
32 for ∀p ∈ Procs(S) \ (Procs(S_new) ∪ {id}) do
33     S_new := S_new ∪ {(p, name(p, S))}
34 end
35 return S_new
```

Figure 3: $(b-1)$-Bounded k-Adaptive Renaming

Lines 6-10 describe how p sets j. If there is any register in p's preceeding view with *writer* component equal to p but with content different from $(S_p, p, name_p)$ then p writes to register $R[j]$ where j is the biggest index amongst these registers. Otherwise it writes to some register whose content is different than $(S_p, p, name_p)$. Process p continues until either in some scan, all registers contain the same information that p has written or $|S_p|$ is larger than or equal $b-1$. In the first case p returns $name_p$ and in the second case it returns $b+p-1$.

Proof of Correctness

Overview of proof Once a process p terminates with name $n_p \leq b-1$, the pair (*writer*, *proposal*) of every register is equal to (p, n_p). The core idea is that after p terminates, every register that is overwritten with the wrong name for p or no name for p, has a distinct *writer* component. Therefore, if a subsequent scan by another process, say q, does not include the correct name for p, the set of processes in that scan is large and q terminates with a name larger than or equal b. If the set of processes in the scan is not large, then there is some writer that is in the *writer* component of at least 2 registers. In that case, we prove that for any such pair of registers with the same writer, the correct name for p is in the register with the larger index. In this way, the algorithm ensures that process q keeps (p, n_p) in its naming set, and discards incorrect names for p.

For our proof, we use the notion and terminology for register configuration, consistent configuration, index(op) of operation op, interval and content of register R in configuration C, $\nu_C(R)$, as defined in Section 5. Let p be a process that has terminated and returned name n_p. Define $last_p$ to be the last scan by p. For any register configuration D following register configuration $C_{\text{index}(last_p)}$, define a set of registers $\mathcal{F}_p(D) = \{R \in \mathcal{R} \mid (\nu_D(R).writer \neq p) \wedge ((p, n_p) \notin \nu_D(R).set)\}$ and a set of processes $O_p(D) = \bigcup_{R \in \mathcal{F}_p(D)} \{\nu_D(R).writer\}$.

Lemma 6.1. *Let E be any execution starting in the initial configuration and ending in configuration C. If there are two integers i and j such that $i < j$, $\nu_C(R[i]).writer = \nu_C(R[j]).writer = p$ and $\nu_C(R[i]) \neq \nu_C(R[j])$, then the last write to $R[i]$ happens before the last write to $R[j]$ and both are by the same process.*

Proof. By Line 2, the *writer* segment of each register indicates the id of the process which writes that value. Hence, $\nu_C(R[i])$ and $\nu_C(R[j])$ are both written by the same process p. Let w_i and w_j be the most recent writes to $R[i]$ and $R[j]$ preceding C, respectively. Thus, value of w_i (respectively w_j) is $\nu_C(R[i])$ (respectively $\nu_C(R[j])$). By way of contradiction assume that w_j happens before w_i. Let s_i be the most recent scan operation by p before w_i. Hence s_i happens after w_j and before w_i. Since w_j is the most recent write to $R[j]$ preceding s_i, $\nu_{C_{\text{index}(s_i)}}(R[j]) = \nu_C(R[j])$. Let $\widehat{S_p}$ and $\widehat{proposed}_p$ be the value of S_p and $proposed_p$ at w_i respectively. Then $(\widehat{S_p}, p, \widehat{proposed}_p) = \nu_C(R[i])$. Therefore, when p executed Line 6 preceding w_i and after s_i, $\widehat{S_p}$ and $\widehat{proposed}_p$ are values of S_p and $proposed_p$ respectively. Let $\widehat{r_j}$ be the value of $r[j]$ at s_i. Thus $\widehat{r_j} = \nu_{C_{\text{index}(s_i)}}(R[j])$. Furthermore, $\widehat{r_j}$ is the value of $r[j]$, when p executes Line 6 after s_i and preceding w_i. Therefore at the execution of Line 6 after s_i and preceding w_i, $(S_p, p, proposed_p) = (\widehat{S_p}, p, \widehat{proposed}_p) = \nu_C(R[i]) \neq \nu_C(R[j]) = \nu_{C_{\text{index}(s_i)}}(R[j]) = \widehat{r_j} = r[j]$ and $r[j].id = \widehat{r_j}.id = p$. Thus, Line 6 evaluates to true. Since $j > i$, by Line 7, p does not write into $R[i]$ before writing into $R[j]$.

□

Informally, Lemma 6.2 says that every register that contains an incorrect name for p after a consistent configuration containing the correct name for p has a distinct *writer* component.

Lemma 6.2. *Consider an execution E in which process p's* `getName()` *call returns name $n_p \leq b-1$. Then for any register configuration C_e where $e \geq \text{index}(last_p)$,*

i) $|\mathcal{F}_p(C_e)| = |O_p(C_e)|$*;*

ii) $\forall q \in O_p(C_e)$*, q performs a write in Interval*$[\text{index}(last_p), e]$*; and*

iii) for any write operation o by any process q during Interval$[\text{index}(last_p), e]$*, let v be the value of o. If o is not q's first write during Interval*$[\text{index}(last_p), e]$*, then $(p, n_p) \in v.set$.*

Proof. Let $C_{\text{index}(last_p)}, C_{\text{index}(last_p)+1}, \ldots$ be the sequence of all register configurations starting at $C_{\text{index}(last_p)}$. We prove the lemma by induction on the indices of this sequence. Let $\widehat{S_p}$ be the value of S_p at $last_p$. For the base case $\ell = \text{index}(last_p)$, since $n_p \leq b-1$, p returns in Line 13. Therefore the condition $r[1] = \cdots = r[b] = (\widehat{S_p}, p, n_p)$

held when p last executed Line 11. Hence, condition $R[1] = \cdots = R[b] = (\widehat{S_p}, p, n_p)$ held at $C_{\text{index}(last_p)}$. Therefore the induction hypothesis (i) and (ii) hold for the base case $\ell = \text{index}(last_p)$ because $\mathcal{F}_p(C_{\text{index}(last_p)}) = O_p(C_{\text{index}(last_p)}) = \emptyset$. Furthermore, since Interval$[\text{index}(last_p), \text{index}(last_p)]$ contains only one write, (iii) is true for the base case.

Suppose that the lemma holds for $\ell - 1 \geq \text{index}(last_p)$. Let w be the write that changes register configuration $C_{\ell-1}$ to C_ℓ, and let x be the process that performs w. Then clearly $x \neq p$, since p has performed its last write before $C_{\text{index}(last_p)}$. Suppose w writes value $(\widehat{S_x}, x, n_x)$ into register R, and let s be x's scan operation that precedes w if it exists.

Suppose $(p, n_p) \in \widehat{S_x}$. Let $\mathcal{F}_p(C_\ell) = \mathcal{F}_p(C_{\ell-1}) \setminus \{R\}$ and $O_p(C_\ell) = O_p(C_{\ell-1}) \setminus \{\nu_{C_{\ell-1}}(R).writer\}$. If $R \in \mathcal{F}_p(C_{\ell-1})$, then by definition, $\nu_{C_{\ell-1}}(R).writer \in O_p(C_{\ell-1})$ and if $R \notin \mathcal{F}_p(C_{\ell-1})$, then by definition, $\nu_{C_{\ell-1}}(R).writer \notin O_p(C_{\ell-1})$. Since $|\mathcal{F}_p(C_{\ell-1})| = |O_p(C_{\ell-1})|$, $|\mathcal{F}_p(C_\ell)| = |O_p(C_\ell)|$. Therefore (i) is true. Since $O_p(C_\ell) \subseteq O_p(C_{\ell-1})$, (ii) holds. Since $(p, n_p) \in \widehat{S_x}$, (iii) is true.

Now consider the case $(p, n_p) \notin \widehat{S_x}$. We first show that $C_{\text{index}(s)}$ precedes $C_{\text{index}(last_p)}$ in Γ_E or w is the first write by x. Suppose, for the purpose of contradiction, that $\text{index}(last_p) \leq \text{index}(s) \leq \ell - 1$. First consider the case that there exists an i such that $\nu_{C_{\text{index}(s)}}(R[i]).writer = p$. Since at $C_{\text{index}(last_p)}$ all registers contain $(\widehat{S_p}, p, n_p)$ and p does not write after $last_p$, $\nu_{C_{\text{index}(s)}}(R[i]).proposal = n_p$. Hence by Line 21, $(p, n_p) \in \widehat{S_x}$. Otherwise suppose that in $C_{\text{index}(s)}$, there are at least two distinct registers whose *writer* are the same process and not p. Then, choose any indices i, j such that $i < j$ and $\nu_{C_{\text{index}(s)}}(R[i]).writer = \nu_{C_{\text{index}(s)}}(R[j]).writer = u \neq p$. Let w_1 and w_2 be the most recent writes to $R[i]$ and $R[j]$ preceding $C_{\text{index}(s)}$. Hence w_1 has value $\nu_{C_{\text{index}(s)}}(R[i])$ and w_2 has value $\nu_{C_{\text{index}(s)}}(R[j])$ and they both are performed by process u. Furthermore, since at $C_{\text{index}(last_p)}$ all registers contain $(\widehat{S_p}, p, n_p)$, w_1 and w_2 occur in Interval$[\text{index}(last_p), \text{index}(s)]$. Suppose $\nu_{C_{\text{index}(s)}}(R[i]) \neq \nu_{C_{\text{index}(s)}}(R[j])$, then by Lemma 6.1, w_1 precedes w_2 in E and by the induction hypothesis (iii), $(p, n_p) \in \nu_{C_{\text{index}(s)}}(R[j]).set$. Otherwise suppose $\nu_{C_{\text{index}(s)}}(R[i]) = \nu_{C_{\text{index}(s)}}(R[j])$ then again by induction hypothesis (iii), $(p, n_p) \in \nu_{C_{\text{index}(s)}}(R[j]).set$. In either case, when x performs Line 24 after s and preceding w, this line evaluates to true. Hence by Line 25, $(p, n_p) \in \widehat{S_x}$. Finally, if $\forall i, j$, $1 \leq i, j \leq b$ and $i \neq j$, $\nu_{C_{\text{index}(s)}}(R[i]).writer \neq \nu_{C_{\text{index}(s)}}(R[j]).writer$, then by induction hypothesis (i), $|\mathcal{F}_p(C_{\text{index}(s)})| = |O_p(C_{\text{index}(s)})| = b$. Therefore, by the for-loop (Lines 18-22), $|\widehat{S_x}| + 1 \geq |O_p(C_{\text{index}(s)})| = b$. Hence, by Line **??**, the presumed write w by x cannot happen. Thus, in all cases, we have established that if $(p, n_p) \notin \widehat{S_x}$ then $C_{\text{index}(s)}$ precedes $C_{\text{index}(last_p)}$ in Γ_E or w is the first write by x.

Consequently, any write by x before w happens before $C_{\text{index}(last_p)}$. On the other hand, by the induction hypothesis, for all $q \in O_p(C_{\ell-1})$, q performs a write during Interval$[\text{index}(last_p), \ell - 1]$ implying x is not in $O_p(C_{\ell-1})$. Thus by defining $\mathcal{F}_p(C_\ell) = \mathcal{F}_p(C_{\ell-1}) \cup \{R\}$ and $O_p(C_\ell) = O_p(C_{\ell-1}) \setminus \{\nu_{C_{\ell-1}}(R).writer\} \cup \{x\}$, the induction hypothesis (i) and (ii) hold for ℓ. Since the most recent operation before w by x happens before $C_{\text{index}(last_p)}$, x performs only one write operation in Interval$[\text{index}(last_p), \ell]$. Therefore, (iii) holds for ℓ.

□

Lemma 6.3. *Let p and q be two distinct processes that have terminated in execution E and returned names n_p and n_q respectively. Suppose that $C_{\text{index}(last_p)}$ precedes $C_{\text{index}(last_q)}$ in Γ_E. If $n_p, n_q \leq b - 1$, then $|\mathcal{F}_p(C_{\text{index}(last_q)})| = 0$.*

Proof. Since $n_q \leq b - 1$, q returns in Line 13. Hence $C_{\text{index}(last_q)}$ is consistent with content (S_q, q, n_q). Therefore, $|\mathcal{F}_p(C_{\text{index}(last_q)})| \in \{0, b\}$. By Lemma 6.2, $|\mathcal{F}_p(C_{\text{index}(last_q)})| = |O_p(C_{\text{index}(last_q)})|$. Since in $C_{\text{index}(last_q)}$, $R.writer = q$ for all $R \in \mathcal{R}$, $|O_p(C_{\text{index}(last_q)})| \leq 1$, and thus $|\mathcal{F}_p(C_{\text{index}(last_q)})| \leq 1$. Therefore $|\mathcal{F}_p(C_{\text{index}(last_q)})| = 0$.

□

Lemma 6.4. *The names returned by any two distinct processes are distinct.*

Proof. For any two distinct processes p and q, let n_p and n_q be the names returned by p and q, respectively. Let $\widehat{S_p}$ (respectively, $\widehat{S_q}$) be the value of S_p (respectively, S_q) when p (respectively, q) executes Line 12. If $|\widehat{S_p}|, |\widehat{S_q}| \geq b - 1$, the names returned by p and q in Line 15 are distinct because $p \neq q$.

Consider the case $|\widehat{S_p}| \leq b - 2$ and $|\widehat{S_q}| \geq b - 1$. Process p returns n_p in Line 13. Since $|\text{Names}(S_p)| \leq |S_p|$, by Line 5, n_p must be smaller than or equal to $b - 1$. Furthermore the name returned by q in Line 15 is larger than or equal b. The case $|\widehat{S_q}| \leq b - 2$ and $|\widehat{S_p}| \geq b - 1$ is true by symmetry. Consider the case $|\widehat{S_p}|, |\widehat{S_q}| \leq b - 2$ implying $n_p, n_q \leq b - 1$. Without loss of generality assume that $C_{\text{index}(last_p)}$ precedes $C_{\text{index}(last_q)}$ in Γ_E. By Lemma 6.3,

$|\mathcal{F}_p(C_{\text{index}(last_q)})| = 0$. Since $\forall R \in \mathcal{R}$, $\nu_{C_{\text{index}(last_q)}}(R).writer = q \neq p$, $(p, n_p) \in \nu_{C_{\text{index}(last_q)}}(R).set$. Thus by Line 23, $(p, n_p) \in \widehat{S_q}$. Therefore by Line 5, $proposed_q \neq n_p$.

□

Observation 6.5. *Let $\widehat{S_p}$ be the value of S_p created by* `Update` *in Line 5 following p's scan operation $scan_p$ in Line 3. Then $\forall q \in \text{Procs}(\widehat{S_p})$, q performs at least one write before $scan_p$.*

Lemma 6.6. *Let k be the number of participating processes during process p's* `getName()`*. Then, the name returned by p, is in the range $\{1, \ldots, k\}$, if $k \leq b-1$ and in the range $\{1, \ldots, n+b-1\}$, if $k \geq b$.*

Proof. Let $\widehat{S_p}$ be the value of S_p when p executes Line 4 for the last time and n_p be the name returned by p. By Observation 6.5, $\forall q \in \text{Procs}(\widehat{S_p})$, q performs at least one write before p returns. Thus, $\forall q \in \text{Procs}(\widehat{S_p})$, q is a participating process. Hence, $|\widehat{S_p}| + 1 \leq k$.

If $k \leq b-1$, process p returns in Line 13. By definition, $|\text{Names}(\widehat{S_p})| \leq |\widehat{S_p}| \leq k-1$. Therefore by Line 5, $n_p \leq |\text{Names}(\widehat{S_p})| + 1 \leq k$.

If $k \geq b$, then p returns either in Line 13 or in Line 15. Therefore the name is in the range $\{1, \ldots, b+n-1\}$.

□

Theorem 6.7. *For any $b \geq 2$ there is an obstruction-free $(b-1)$-bounded k-adaptive renaming algorithm implemented from $b+1$ bounded registers such that when $k \geq b$ the returned names are in the range $\{1, \ldots, n+b-1\}$.*

Proof. There is an obstruction-free implementation of b-component snapshot objects from $b+1$ bounded registers [22]. Since our algorithm in Fig. 3 is deterministic we can replace the atomic scan registers with a linearizable scan. By Lemma 6.6 and Lemma 6.4, the algorithm solves $(b-1)$-bounded k-adaptive renaming. Thus, it suffices to prove that the algorithm is obstruction-free.

If p runs alone then the value of S_p computed in Line 4 and $proposed_p$ computed in Line 5 remain the same. Therefore after b write operations all registers contain $(S_p, p, proposed_p)$. Therefore, in the b-th iteration of the repeat-until loop (Line 11) evaluates to true and p stops.

□

Let $f : \{1, \ldots, n\} \to \mathbb{N}$ be a non-decreasing function where, $\forall k \in \{1, \ldots, n\}$, $f(k) \geq k$ and $f(1) \leq n-1$. Let $d' = \min\{n, x \mid f(x) \geq 2n\}$. Hence, $d' \leq n$. Setting $b = d'$, we have an obstruction-free one-shot d'-bounded k-adaptive renaming algorithm from $d'+1$ registers. This implies that the algorithm returns names in the range $\{1, \ldots, k\}$ when $k \leq d'-1$, and returns names in the range $\{1, \ldots, n+d'-1\}$ when $k \geq d'$. Note that $\forall k \in \{1, \ldots, d'-1\}$, $k \leq f(k)$. Furthermore, when $k \geq d'$, $n+d'-1 < 2n \leq f(d')$. Hence, $\forall k \in \{1, \ldots, n\}$, the algorithm returns names in the range $\{1, \ldots, f(k)\}$.

Corollary 6.1. *There is an obstruction-free one-shot f-adaptive renaming algorithm implemented from $\min\{n, x \mid f(x) \geq 2n\} + 1$ bounded registers.*

7 Observations and Open Problems

Let $f : \{1, \ldots, n\} \to \mathbb{N}$ be a non-decreasing function where, $\forall k \in \{1, \ldots, n\}$, $f(k) \geq k$ and $f(1) \leq n-1$. Let $\doteq \max\{x \mid f(x) \leq n-1\}$. We proved a lower bound of $\doteq 1$ for non-deterministic solo-terminating long-lived f-adaptive renaming. Furthermore, for any integer constant $0 \leq c \leq n$, we showed a lower bound of $\lfloor \frac{2(n-c)}{c+2} \rfloor$ for non-deterministic solo-terminating one-shot $(k+c)$-adaptive renaming. This implies a tight space bound of n for both one-shot and long-lived tight renaming. We also presented an obstruction-free one-shot f-adaptive algorithm from $\min\{n, x \mid f(x) \geq 2n\} + 1$ registers.

An obvious solution for any obstruction-free long-lived or one-shot f-adaptive renaming is as follows. A set $Q \subseteq \mathcal{P}$ of $\lfloor f(1) \rfloor - 1$ processes always return names in the range $\{1, \ldots, \max(\lfloor f(1) \rfloor - 1, 1)\}$ without taking any steps. In any $(\mathcal{P} \backslash Q)$-solo execution, process in $\mathcal{P} \backslash Q$ using universal construction, get names in the range $\{\lfloor f(1) \rfloor, \ldots, k + \lfloor f(1) \rfloor - 1\}$. Universal construction for $|\mathcal{P} \backslash Q|$ processes requires $|\mathcal{P}| - |Q| = n - \lfloor f(1) \rfloor + 1$ registers. Observe that

this is a tight upper bound for obstruction-free long-lived $(k+c)$-adaptive renaming. One of the most noticeable open problems is whether implementing one-shot f-adaptive renaming requires asymptotically less space than long-lived f-adaptive renaming.

We designed a wait-free one-shot $(b-1)$-bounded $(k(k+1)/2)$-adaptive renaming algorithm from b bounded registers, and established that this algorithm has a polynomial step complexity. It appears that if we modify the `newScan()` function of our algorithm, so that each process returns when the set of all processes know to it grows even by one, the step complexity would reduce considerably. However this change would require much more elaborate and challenging proofs because the values returned by `newScan()` would not be equivalent to values returned by a linearizable scan.

For some systems, it seems reasonable to have the register space, as well as the name space, adapt to the actual number of participants. The one-shot lower bound can also be modified to express the actual register use as a function of k. On the other hand, the one-shot algorithms in this paper require a fixed number of registers regardless of the number of participants.

References

[1] Y. Afek and M. Merritt. Fast, wait-free (2k-1)-renaming. In *Proc. of 18th PODC*, pages 105–112. Journal of the ACM, 1999.

[2] D. Alistarh, J. Aspnes, K. Censor-Hillel, S. Gilbert, and M. Zadimoghaddam. Optimal-time adaptive strong renaming, with applications to counting. In *Proc. of 30th PODC*, pages 239–248, 2011.

[3] D. Alistarh, J. Aspnes, S. Gilbert, and R. Guerraoui. The complexity of renaming. In *Proc. of 52nd FOCS*, pages 718–727, 2011.

[4] D. Alistarh, H. Attiya, S. Gilbert, A. Giurgiu, and R. Guerraoui. Fast randomized test-and-set and renaming. In *Proc. of 24th DISC*, pages 94–108, 2010.

[5] H. Attiya, A. Bar-Noy, D. Dolev, D. Peleg, and R. Reischuk. Renaming in an asynchronous environment. *Journal of the ACM*, 37(3):524–548, 1990.

[6] H. Attiya, F. Fich, and Y. Kaplan. Lower bounds for adaptive collect and related objects. In *Proc. of 23rd PODC*, pages 60–69, 2004.

[7] H. Attiya and A. Paz. Counting-based impossibility proofs for renaming and set agreement. In *Proc. of 26th DISC*, pages 356–370, 2012.

[8] H. Attiya and J. Welch. *Distributed Computing: Fundamentals, Simulations and Advanced Topics*. John Wiley Interscience, 2004.

[9] E. Borowsky and E. Gafni. Generalized flp impossibility result for t-resilient asynchronous computations. In *Proc. of 25th ACM STOC*, pages 91–100, 1993.

[10] A. Brodsky, F. Ellen, and P. Woelfel. Fully-adaptive algorithms for long-lived renaming. *Journal of Distributed Computing*, 24(2):119–134, 2011.

[11] J. Burns and N. Lynch. Bounds on shared memory for mutual exclusion. *Journal of Information and Computation*, 107(2):171–184, 1993.

[12] J. Burns and G. Peterson. The ambiguity of choosing. In *Proc. of 8th PODC*, PODC '89, pages 145–157. Journal of the ACM, 1989.

[13] A. Castañeda and S. Rajsbaum. New combinatorial topology upper and lower bounds for renaming. In *Proc. of 27th PODC*, pages 295–304, 2008.

[14] Danny Dolev and Nir Shavit. Bounded concurrent time-stamping. *SIAM Journal on Computing*, 26(2):418–455, 1997.

[15] Cynthia Dwork and Orli Waarts. Simple and efficient bounded concurrent timestamping and the traceable use abstraction. *Journal of the ACM*, 46(5):633–666, 1999.

[16] W. Eberly, L. Higham, and J. Warpechowska-Gruca. Counting-based impossibility proofs for renaming and set agreement. In *Proc. of 26th DISC*, pages 149–160, 1998.

[17] F. Ellen and P. Woelfel. An optimal implementation of fetch-and-increment. In *Proc. of 27th DISC*, pages 284–298, 2013.

[18] F. Fich, M. Herlihy, and N. Shavit. On the space complexity of randomized synchronization. *Journal of the ACM*, pages 843–862, 1998.

[19] F. Fich and E. Ruppert. Hundreds of impossibility results for distributed computing. *Distributed Computing*, 16(2-3):121–163, 2003.

[20] E. Gafni. Read-write reductions. In *Proceedings of the 8th International Conference on Distributed Computing and Networking*, pages 349–354, 2006.

[21] Rainer Gawlick, Nancy A. Lynch, and Nir Shavit. Concurrent timestamping made simple. In *1st Israel Symposium on Theory of Computing Systems (ISTCS)*, pages 171–183, 1992.

[22] G. Giakkoupis, M. Helmi, L. Higham, and P. Woelfel. An O(sqrt n) space bound for obstruction-free leader election. In *Proc. of 31st DISC*, pages 46–60, 2013.

[23] M. Helmi, L. Higham, E. Pacheco, and P. Woelfel. The space complexity of long-lived and one-shot timestamp implementations. *Journal of the ACM*, 61(1):7–27, 2014.

[24] M. Herlihy and N. Shavit. The topological structure of asynchronous computability. *Journal of the ACM*, 46(6):858–923, 1999.

[25] Amos Israeli and Ming Li. Bounded time-stamps. *Distributed Computing*, 6(4):205–209, 1993.

[26] Amos Israeli and Meir Pinhasov. A concurrent time-stamp scheme which is linear in time and space. In *Distributed Algorithms, 6th International Workshop (WDAG)*, pages 95–109, 1992.

[27] N. Lynch. *Distributed Algorithms*. Morgan Kaufmann, 1996.

[28] M. Moir and J. Garay. Fast, long-lived renaming improved and simplified. In *Procedeeding of the 10th International Workshop on Distributed Algorithms (WDAG)*, pages 287–303, 1996.

[29] A. Panconesi, M. Papatriantafilou, P. Tsigas, and P. Vitnyi. Randomized naming using wait-free shared variables. *Journal of Distributed Computing*, 11:113–124, 1998.